Networking the world's business data™

McDATA PRODUCTS

McDATA®
Products in a SAN Environment
Planning Manual

P/N 620-000124-510
REV A

Simplifying Storage Network Management

McDATA Corporation
380 Interlocken Crescent Broomfield, CO 80021-3464
Corporate Headquarters: 800-545-5773
Sales E-mail: sales@mcdata.com *Web:* www.mcdata.com

CERTIFIED QUALITY MANAGEMENT SYSTEM ISO 9001:2000
ANSI-RAB QMS ACCREDITED
DNV

Record of Revisions and Updates

Revision	Date	Description
620-000124-000	5/2002	Initial release of the manual.
620-000124-100	9/2002	Revision of the manual to describe the Intrepid 6140 Director, Sphereon 4500 Fabric Switch, and Release 6.3 of the Enterprise Fabric Connectivity Manager application.
620-000124-200	2/2003	Revision of the manual to include additional information and describe Release 7.1 of the Enterprise Fabric Connectivity Manager application.
620-000124-300	8/2003	Revision of the manual to describe the Sphereon 4300 Fabric Switch and Release 7.2 of the Enterprise Fabric Connectivity Manager application.
620-000124-400	12/2003	Revision of the manual to describe Release 8.1 of the Enterprise Fabric Connectivity Manager application.
620-000124-500	2/2005	Revision of the manual to describe the Eclipse 1620 SAN Router, Eclipse 2640 SAN Router, Intrepid 10000 Director, Release 4.6 of the SANvergence Manager application, and Release 8.6 of the Enterprise Fabric Connectivity Manager application.
620-000124-510	7/2005	Revision of the manual to describe the Sphereon 4400 Fabric Switch, Sphereon 4700 Fabric Switch, Release 8.0 of the Enterprise Operating System, and Release 8.7 of the Enterprise Fabric Connectivity Manager application

Copyright © 2002 - 2005 McDATA Corporation. All rights reserved.

Printed July 2005

Seventh Edition

No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written consent of McDATA Corporation.

The information contained in this document is subject to change without notice. McDATA Corporation assumes no responsibility for any errors that may appear.

All computer software programs, including but not limited to microcode, described in this document are furnished under a license, and may be used or copied only in accordance with the terms of such license. McDATA either owns or has the right to license the computer software programs described in this document. McDATA Corporation retains all rights, title and interest in the computer software programs.

McDATA Corporation makes no warranties, expressed or implied, by operation of law or otherwise, relating to this document, the products or the computer software programs described herein. McDATA CORPORATION DISCLAIMS ALL IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE. In no event shall McDATA Corporation be liable for (a) incidental, indirect, special, or consequential damages or (b) any damages whatsoever resulting from the loss of use, data or profits, arising out of this document, even if advised of the possibility of such damages.

Contents

Preface xiii

Chapter 1 Introduction to McDATA Multi-Protocol Products

Product Overview 1-2
- Multi-Protocol Hardware 1-2
- SAN Management Applications 1-5

Directors 1-6
- Director Performance 1-7
- Intrepid 6064 Director 1-8
- Intrepid 6140 Director 1-10
- Intrepid 10000 Director 1-12

Fabric Switches 1-14
- Fabric Switch Performance 1-15
- Sphereon 3232 Fabric Switch 1-15
- Sphereon 4300 Fabric Switch 1-16
- Sphereon 4400 Fabric Switch 1-18
- Sphereon 4500 Fabric Switch 1-19
- Sphereon 4700 Fabric Switch 1-21

SAN Routers 1-22
- SAN Router Performance 1-23
- Eclipse 1620 SAN Router 1-24
- Eclipse 2640 SAN Router 1-26

Product Features 1-28
- Connectivity Features 1-28
- Security Features 1-31
- Serviceability Features 1-32

Chapter 2 Product Management

Product Management....2-2
Out-of-Band Management....2-2
Inband Management....2-5
Management Interface Summary....2-6
Management Server Support....2-7
Management Server Specifications....2-8
Ethernet Hub....2-10
Remote User Workstations....2-10
Product Firmware....2-11
Firmware Services....2-12
Backup and Restore Features....2-13
SAN Management Applications....2-15
SANavigator and EFCM Applications....2-15
SANvergence Manager Application....2-18
EFCM Basic Edition Interface....2-20
Command Line Interface....2-21

Chapter 3 Planning Considerations for Fibre Channel Topologies

Fibre Channel Topologies....3-1
Characteristics of Arbitrated Loop Operation....3-2
Shared Mode Versus Switched Mode....3-3
Public Versus Private Devices....3-6
Public Versus Private Loops....3-8
FL_Port Connectivity....3-10
Planning for Private Arbitrated Loop Connectivity....3-10
Planning for Fabric-Attached Loop Connectivity....3-11
Connecting FC-AL Devices to a Switched Fabric....3-11
Server Consolidation....3-12
Tape Device Consolidation....3-13
Fabric Topologies....3-14
Mesh Fabric....3-14
Core-to-Edge Fabric....3-16
SAN Islands....3-18
Planning for Multiswitch Fabric Support....3-18
Fabric Topology Limits....3-19
Factors to Consider When Implementing a Fabric Topology....3-20
General Fabric Design Considerations....3-29
Fabric Initialization....3-29
Fabric Performance....3-30
Fabric Availability....3-37

Fabric Scalability 3-39
Obtaining Professional Services 3-40
Mixed Fabric Design Considerations 3-40
FCP and FICON in a Single Fabric 3-41
Multiple Data Transmission Speeds in a Single Fabric 3-51
FICON Cascading 3-52
High-Integrity Fabrics 3-53
Minimum Requirements 3-53
FICON Cascading Best Practices 3-54

Chapter 4 Implementing SAN Internetworking Solutions

SAN Island Consolidation 4-2
Flexible Partitioning Technology 4-4
SAN Routing 4-8
Implementing BC/DR Solutions 4-36
Extended-Distance Operational Modes 4-37
SAN Extension Transport Technologies 4-39
Distance Extension Through BB_Credit 4-49
Intelligent Port Speed 4-51
Distance Extension Best Practices 4-55
Consolidating and Integrating iSCSI Servers and Storage 4-58
iSCSI Protocol 4-59
iSCSI Server Consolidation 4-60
iSCSI Storage Consolidation 4-61

Chapter 5 Physical Planning Considerations

Port Connectivity and Fiber-Optic Cabling 5-1
Port Requirements 5-2
SFP Optical Transceivers 5-4
Extended-Distance Ports 5-6
High-Availability Considerations 5-6
Fibre Channel Cables and Connectors 5-7
Routing Fiber-Optic Cables 5-8
Management Server, LAN, and Remote Access Support 5-9
Management Server 5-10
Remote User Workstations 5-11
SNMP Management Workstations 5-13
EFCM Basic Edition Interface 5-14
Security Provisions 5-15
Password Protection 5-15
SANtegrity Authentication 5-16

SANtegrity Binding....5-19
PDCM Arrays....5-20
Preferred Path....5-23
Zoning....5-25
Server and Storage-Level Access Control....5-29
Security Best Practices....5-30
Optional Feature Keys....5-33
Inband Management Access....5-35
Flexport Technology....5-36
SANtegrity Authentication....5-36
SANtegrity Binding....5-36
OpenTrunking....5-37
Full Volatility....5-38
Full Fabric....5-38
Remote Fabric....5-39
N_Port ID Virtualization....5-39
Element Manager Application....5-39

Chapter 6 Configuration Planning Tasks

Task 1: Prepare a Site Plan....6-2
Task 2: Plan Fibre Channel Cable Routing....6-3
Task 3: Consider Interoperability with Fabric Elements and End Devices....6-4
Task 4: Plan Console Management Support....6-5
Task 5: Plan Ethernet Access....6-7
Task 6: Plan Network Addresses....6-7
Task 7: Plan SNMP Support (Optional)....6-10
Task 8: Plan E-Mail Notification (Optional)....6-11
Task 9: Establish Product and Server Security Measures....6-11
Task 10: Plan Phone Connections....6-12
Task 11: Diagram the Planned Configuration....6-12
Task 12: Assign Port Names and Nicknames....6-13
Rules for Port Names....6-13
Rules for Nicknames....6-14
Task 13: Complete the Planning Worksheet....6-14
Task 14: Plan AC Power....6-28
Task 15: Plan a Multiswitch Fabric (Optional)....6-29
Task 16: Plan Zone Sets for Multiple Products (Optional)....6-30
Task 17: Plan SAN Routing (Optional)....6-31
Task 18: Complete Planning Checklists....6-34

Appendix A Product Specifications

Director, Fabric Switch, and SAN Router Specifications A-1
- Dimensions A-1
- Power Requirements A-3
- Heat Dissipation A-5
- Clearances A-5
- Acoustical Noise and Physical Tolerances A-6
- Storage and Shipping Environment A-6
- Operating Environment A-7

FC-512 Fabricenter Cabinet Specifications A-7
- Dimensions A-8
- Power Requirements A-8
- Clearances A-8
- Cabinet Footprint A-8

Appendix B Firmware Summary

System-Related Differences B-1
Fibre Channel Protocol-Related Differences B-3
Management-Related Differences B-5

Index I-1

Tables

2-1 Out-of-Band and Inband Product Support Summary 2-6

4-1 mSAN Routing Domain 4-18
4-2 mSAN Supported Limits 4-21
4-3 mFCP Versus iFCP 4-28
4-4 Transport Technology Comparison 4-48

5-1 Cable Type and Transmission Rate versus Distance and Link Budget 5-5
5-2 Types of User Rights 5-15

6-1 Configuration Planning Tasks 6-1
6-2 Physical Planning and Hardware Installation Tasks 6-35
6-3 Operational Setup Tasks 6-36

B-1 E/OSc versus E/OSn and E/OSi - System-Related Differences B-1
B-2 E/OSc versus E/OSn and E/OSi - Fibre Channel Protocol-Related Differences B-3
B-3 E/OSc versus E/OSn and E/OSi - Management-Related Differences . B-5

Figures

1-1 Cabinet-Mount McDATA Products 1-4
1-2 Intrepid 6064 Director 1-9
1-3 Intrepid 6140 Director 1-11
1-4 Intrepid 10000 Director 1-12
1-5 Sphereon 3232 Fabric Switch 1-15
1-6 Sphereon 4300 Fabric Switch 1-16
1-7 Sphereon 4400 Fabric Switch 1-18
1-8 Sphereon 4500 Fabric Switch 1-19
1-9 Sphereon 4700 Fabric Switch 1-21
1-10 Eclipse 1620 SAN Router 1-25
1-11 Eclipse 2640 SAN Router 1-26

2-1 Out-of-Band Product Management 2-4
2-2 Inband Product Management 2-6
2-3 Management Server 2-7
2-4 24-Port Ethernet Hub 2-10
2-5 Hardware View 2-17
2-6 Main Window (SANvergence Manager) 2-18
2-7 Device Window (Element Manager) 2-19
2-8 Hardware View (EFCM Basic Edition Interface) 2-20

3-1 Shared Mode Operation and Logical Equivalent 3-3
3-2 Switched Mode Operation and Logical Equivalent 3-4
3-3 Public Device Connectivity 3-6
3-4 Private Device Connectivity 3-7
3-5 Public Loop Connectivity 3-9
3-6 Private Loop Connectivity 3-9
3-7 Server Consolidation 3-13
3-8 Tape Drive Consolidation 3-14

3-9 Full Mesh Fabric 3-15
3-10 2-by-14 Core-to-Edge Fabric 3-17
3-11 Example Multiswitch Fabric 3-19
3-12 ISL Oversubscription 3-32
3-13 Device Locality 3-34
3-14 Device Fan-Out Ratio 3-35
3-15 Fabric Performance Tuning 3-36
3-16 Redundant Fabrics 3-39
3-17 Intrepid 6140 Port Numbers and Logical Port Addresses (Front) 3-43
3-18 Intrepid 6140 Port Numbers and Logical Port Addresses (Rear) 3-44

4-1 Intrepid 10000 Director FlexPar Functionality 4-5
4-2 SAN Routing Hierarchy 4-9
4-3 SAN Routing Concepts 4-10
4-4 SAN Routing - Physical Connectivity 4-11
4-5 SAN Routing - Logical Connectivity 4-13
4-6 iFCP WAN Extension 4-24
4-7 Inter-FlexPar Routing 4-29
4-8 Dark Fiber Extended-Distance Connectivity 4-40
4-9 WDM Extended-Distance Connectivity 4-41
4-10 SONET Extended-Distance Connectivity 4-43
4-11 SoIP Extended-Distance Connectivity 4-46
4-12 SAN Extension Technology Comparison 4-47
4-13 WAN Link Performance (No Rate Limiting) 4-51
4-14 WAN Link Performance (Rate Limiting Enabled) 4-52
4-15 iSCSI Server Consolidation 4-60
4-16 iSCSI Storage Consolidation 4-61

5-1 SFP Transceiver and LC Duplex Connector 5-8
5-2 Typical Network Configuration (One Ethernet Connection) 5-12
5-3 Typical Network Configuration (Two Ethernet Connections) 5-13
5-4 Configure Allow / Prohibit Matrix - Active Dialog Box 5-21
5-5 PDCM Array - Example Problem 5-22
5-6 Preferred Path Configuration 5-23
5-7 Director Zoning 5-25
5-8 OpenTrunking 5-37

A-1 Fabricenter Cabinet Footprint A-9

Preface

This publication is part of a documentation suite that supports McDATA® multi-protocol switching and routing products, including the:

- Intrepid® 6064 Director.
- Intrepid 6140 Director.
- Intrepid 10000 Director.
- Sphereon™ 3232 Fabric Switch.
- Sphereon 4300 Fabric Switch.
- Sphereon 4400 Fabric Switch.
- Sphereon 4500 Fabric Switch.
- Sphereon 4700 Fabric Switch.
- Eclipse™ 1620 SAN Router.
- Eclipse 2640 SAN Router.

Who Should Use This Manual

Use this publication if you are planning to acquire and install one or more fabric switching or routing products. The publication describes product features, hardware, software, planning considerations, and planning tasks. The information provided is intended for use by configuration and installation planners; however information is also provided for system administrators, customer engineers, and project managers.

Organization of This Manual

This publication includes six chapters and two appendices organized as follows:

Chapter 1, Introduction to McDATA Multi-Protocol Products - This chapter provides an overview of McDATA multi-protocol products, and describes product performance and connectivity, security, and serviceability features.

Chapter 2, Product Management - This chapter describes out-of-band and inband product management; the management server; product firmware; backup and restore features; and software. Overviews of the graphical user interfaces (GUIs) and command line interface (CLI) are included.

Chapter 3, Planning Considerations for Fibre Channel Topologies - This chapter describes Fibre Channel topologies (including arbitrated loop and fabric); multiswitch fabric topologies and storage area networks (SANs); multiswitch fabric support; general, large, and mixed fabric design considerations; and fibre connection (FICON) cascading.

Chapter 4, Implementing SAN Internetworking Solutions - This chapter describes SAN island consolidation; implementing business continuance and disaster recovery (BC / DR) solutions; and consolidating and integrating Internet small computer systems interface (iSCSI) servers and storage.

Chapter 5, Physical Planning Considerations - This chapter describes physical factors to consider when planning a Fibre Channel SAN configuration. Factors include port connectivity and fiber-optic cabling; management server support; local area network (LAN) and remote access support; inband management access; and security and zoning support.

Chapter 6, Configuration Planning Tasks - This chapter describes planning tasks to be performed prior to installing a director, fabric switch, or SAN router. Tasks include physical site planning, connectivity and management access, and facility support. A worksheet that lists product port connections is included. Checklists that summarize planning and installation activities are also included.

Appendix A, Product Specifications - This appendix lists specifications for directors, fabric switches, and SAN routers.

Appendix B, Firmware Summary - This appendix summarizes differences and similarities between firmware versions that support directors, fabric switches, and SAN routers.

An *Index* is also provided.

Related Publications

Other publications that provide additional information about McDATA products include:

- **Intrepid 6064 and 6140 Directors:**
  - — *McDATA Intrepid 6064 and 6140 Directors Element Manager User Manual* (620-000172).
  - — *McDATA Intrepid 6064 Director Installation and Service Manual* (620-000108).
  - — *McDATA Intrepid 6140 Director Installation and Service Manual* (620-000157).
- **Intrepid 10000 Director:**
  - — *McDATA Intrepid 10000 Director Element Manager User Manual* (620-000227).
  - — *McDATA Intrepid 10000 Director Installation and Service Manual* (620-000225).
- **Sphereon 3232 Fabric Switch:**
  - — *McDATA Sphereon 3232 Fabric Switch Element Manager User Manual* (620-000173).
  - — *McDATA Sphereon 3232 Fabric Switch Installation and Service Manual* (620-000155).
- **Sphereon 4300 Fabric Switch:**
  - — *McDATA Sphereon 4300 Fabric Switch Installation and Service Manual* (620-000171).
- **Sphereon 4400 Fabric Switch:**
  - — *McDATA Sphereon 4400 Fabric Switch Element Manager User Manual* (620-000241).
  - — *McDATA Sphereon 4400 Fabric Switch Installation and Service Manual* (620-000238).
- **Sphereon 4500 Fabric Switch:**
  - — *McDATA Sphereon 4500 Fabric Switch Element Manager User Manual* (620-000175).
  - — *McDATA Sphereon 4500 Fabric Switch Installation and Service Manual* (620-000159).

- **Sphereon 4700 Fabric Switch:**
  - *McDATA Sphereon 4700 Fabric Switch Element Manager User Manual* (620-000242).
  - *McDATA Sphereon 4700 Fabric Switch Installation and Service Manual* (620-000239).
- **Eclipse 1620 SAN Router:**
  - *McDATA Eclipse 1620 SAN Router Administration and Configuration Manual* (620-000206).
  - *McDATA Eclipse 1620 SAN Router Installation and Service Manual* (620-000205).
- **Eclipse 2640 SAN Router:**
  - *McDATA Eclipse 2640 SAN Router Administration and Configuration Manual* (620-000203).
  - *McDATA Eclipse 2640 SAN Router Installation and Service Manual* (620-000202).
- **General Support Publications**:
  - *SANavigator Software Release 4.2 User Manual* (621-000013).
  - *EFC Manager Software Release 8.7 User Manual* (620-000170).
  - *McDATA SANvergence Manager User Manual* (620-000189).
  - *McDATA EFCM Basic Edition User Manual* (620-000240).
  - *McDATA E/OSc SNMP Agent User Manual* (620-000168).
  - *McDATA E/OSn SNMP Support Manual* (620-000226).
  - *McDATA E/OSi SNMP Support Manual* (620-000228).
  - *McDATA E/OSc Command Line Interface User Manual* (620-000134).
  - *McDATA E/OSn Command Line Interface User Manual* (620-000211).
  - *McDATA E/OSi Command Line Interface User Manual* (620-000207).
  - *McDATA SDK C-FCSWAPI User Manual* (620-000149).

— *McDATA EFCM Lite Installation Instructions* (958-000171).

— *McDATA FC-512 Fabricenter Equipment Cabinet Installation and Service Manual* (620-000100).

Ordering Printed Manuals

To order a printed copy of this publication, submit a purchase order as described in *Ordering McDATA Documentation Instructions* at http://www.mcdata.com. To obtain documentation CD-ROMs, contact your McDATA sales representative.

Where to Get Help

For technical support, contact the McDATA Solution Center. The center provides a single point of contact for assistance and is staffed 24 hours a day, seven days a week, including holidays. Contact the center at the phone number, fax number, or e-mail address listed below. Please have the product serial number (printed on the service label) available.

Phone: (800) 752-4572 or (720) 558-3910

Fax: (720) 558-3851

E-mail: support@mcdata.com

Trademarks

The following terms, indicated by a registered trademark symbol (®) or trademark symbol (™) on first use in this publication, are trademarks of McDATA Corporation or SANavigator, Inc. in the United States or other countries or both:

Registered Trademarks	Trademarks
Fabricenter®	Eclipse™
HotCAT®	EON™
Intrepid®	FlexPar™
McDATA®	nScale™
Multi-Capable Storage Network Solutions®	OPENconnectors™
Networking the World's Business Data®	Sphereon™
OPENready®	
SANavigator®	
SANtegrity®	
SANvergence®	

All other trademarked terms, indicated by a registered trademark symbol (®) or trademark symbol (™) on first use in this publication, are trademarks of their respective owners in the United States or other countries or both.

Forwarding Publication Comments

Please send comments to the McDATA Solution Center by telephone, fax, or e-mail. The numbers and e-mail address are listed above. Please identify the page numbers and details.

Laser Compliance Statement

Product laser transceivers are tested and certified in the United States to conform to Title 21 of the Code of Federal Regulations (CFR), Subchapter J, Parts 1040.10 and 1040.11 for Class 1 laser products. Transceivers are tested and certified to be compliant with International Electrotechnical Commission IEC825-1 and European Norm EN60825-1 and EN60825-2 regulations for Class 1 laser products. Class 1 laser products are not considered hazardous. The transceivers are designed to prevent human access to laser radiation above a Class 1 level during normal operation or prescribed maintenance conditions.

Federal Communications Commission (FCC) Statement

Products generate, use, and can radiate radio frequency energy, and if not installed and used in accordance with instructions provided, may cause interference to radio communications. Products are tested and found to comply with the limits for Class A and Class B computing devices pursuant to Subpart B of Part 15 of the FCC Rules, which are designed to provide reasonable protection against such interference in a residential environment. Any modification or change made to a product without explicit approval from McDATA, by means of a written endorsement or through published literature, invalidates the service contract and voids the warranty agreement with McDATA.

Canadian EMC Statements

The statements below indicate product compliance with Interference Causing Equipment Standard (ICES) and Norme sur le Matériel Brouiller (NMB) electromagnetic compatibility (EMC) requirements as set forth in ICES / NMB-003, Issue 4.

- This Class A or Class B digital apparatus complies with Canadian ICES-003.
- Cet appareil numérique de la classe A et classe B est conforme à la norme NMB-003 du Canada.

United States and Canada UL Certification

The C-UL-US mark on a product indicates compliance with American National Standards Institute (ANSI) and Standards Council of Canada (SCC) safety requirements as tested, evaluated, and certified by Underwriters Laboratories Inc. (UL) and Underwriters Laboratories of Canada (ULC).

International Safety Conformity Declaration (CB Scheme)

A certification bodies (CB) test report supporting a product indicates safety compliance with the International Electrotechnical Commission (IEC) system for conformity testing and certification of electrical equipment (IECEE) CB scheme.

The CB scheme is a multilateral agreement among participating countries and certification organizations that accepts test reports certifying the safety of electrical and electronic products.

European Union Conformity Declarations and Directives (CE Mark)

The CE mark on a product indicates compliance with the following regulatory requirements as set forth by European Norms (ENs) and relevant international standards for commercial and light industrial information technology equipment (ITE):

- **EN55022: 1998 -** ITE-generic radio frequency interference (RFI) emission standard for domestic, commercial, and light industrial environments, including electrical business equipment.
- **EN55024-1: 1998 -** ITE-generic electromagnetic immunity standard for domestic, commercial, and light industrial environments, including electrical business equipment.
- **EN60950/A11:1997 -** ITE-generic electrical and fire safety standard for domestic, commercial, and light industrial environments, including electrical business equipment.
- **EN61000-3-2:1995 -** ITE-generic harmonic current emissions standard for domestic, commercial, and light industrial environments (equipment with rated current less than or equal to 16 amperes per phase).
- **EN61000-3-3:1995 -** ITE-generic voltage fluctuation and flicker standard (low-voltage power supply systems) for domestic, commercial, and light industrial environments (equipment with rated current less than or equal to 16 amperes per phase).

In addition, the European Union (EU) Council has implemented a series of directives that define product safety standards for member countries. The following directives apply:

- Products conform with all protection requirements of EU directive **89/336/EEC** (Electromagnetic Compatibility Directive) in accordance with the laws of the member countries relating to EMC emissions and immunity.
- Products conform with all protection requirements of EU directive **73/23/EEC** (Low-Voltage Directive) in accordance with the laws of the member countries relating to electrical safety.
- Products conform with all protection requirements of EU directive **93/68/EEC** (Machinery Directive) in accordance with the laws of the member countries relating to safe electrical and mechanical operation of the equipment.

McDATA does not accept responsibility for any failure to satisfy the protection requirements of any of these directives resulting from a non-recommended or non-authorized modification to a product.

European Union EMC and Safety Declaration (N-Mark)

The N-mark on a product indicates compliance with European Union EMC and safety requirements as tested, evaluated, and certified by the Norwegian Board for Testing and Approval of Electrical Equipment (Norges Elektriske Materiellkontroll or NEMKO) laboratory or a NEMKO-authorized laboratory.

Argentina IRAM Certification

The Instituto Argentino de Normalización (IRAM) S-mark on a product indicates compliance with Direccion Nacional de Comercio Interior (DNCI) Resolution Number 92/98, Phase III (for information technology equipment safety). In conjunction with the S-mark is the AR-UL mark, certified by UL de Argentina, S.R.L., and accredited by the Argentine Accreditation Organization (OAA).

Australia and New Zealand C-Tick Mark

The Australia and New Zealand regulatory compliance mark (C-tick mark) on a product indicates compliance with regulatory requirements for safety and EMC (for information technology equipment) as set forth by the Australian Communications Authority (ACA) and the Radio Spectrum Management Group (RSM) of New Zealand.

People's Republic of China CCC Mark

The China Compulsory Certification mark (CCC mark) on a product indicates compliance with People's Republic of China regulatory requirements for safety and EMC (for information technology equipment) as set forth by the National Regulatory Commission for Certification and Accreditation.

Chinese National Standards Statement

The Chinese National Standards (CNS) statement below indicates product compliance with Taiwanese Bureau of Standards, Metrology, and Inspection (BSMI) regulatory requirements. The statement indicates a product is a Class A or Class B product, and in a domestic environment may cause radio interference, in which case the user is required to take corrective actions.

這是乙類的資訊產品，在居住環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

German TÜV GS Mark

The German regulatory compliance mark (TÜV GS Mark) on a product indicates compliance with the German Safety of Equipment Act as tested by the Technical Inspection Association (Technischer Überwachungsverein or TÜV), and accredited by the Central Office of Safety of the German Länder (Zentralstelle der Länder für Sicherheit or ZLS).

Japanese VCCI Statement

The Voluntary Control Council for Interference (VCCI) statement below applies to information technology equipment, and indicates product compliance with Japanese regulatory requirements. The statement indicates a product is a Class A or Class B product, and in a domestic environment may cause radio interference, in which case the user is required to take corrective actions.

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスB情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean MIC Mark

The Korean Ministry of Information and Communications mark (MIC mark) on a product indicates compliance with regulatory requirements for safety and EMC (for information technology equipment) as authorized and certified by the Korean Radio Research Institute (RRI).

Mexican NOM Mark

The Official Mexican Standard (Normas Oficiales Mexicanas or NOM) mark on a product indicates compliance with regulatory requirements for safety (for information technology equipment) as authorized and accredited by the National System of Accreditation of Testing Laboratories (Sistema Nacional de Acreditamieno de Laboratorios de Pruebas or SINALP).

Russian GOST Certification

The Russian Gosudarstvennyi Standart (GOST) mark on a product indicates compliance with regulatory requirements for safety and EMC (for information technology equipment) as authorized and accredited by the State Committee for Standardization, Metrology and Certification.

Danger and Attention Statements

The following **DANGER** statement appears in this publication and describes safety practices that must be observed while installing or servicing a product. A **DANGER** statement provides essential information or instructions for which disregard or noncompliance may result in death or severe personal injury. The statement appears in English, followed by translations to:

- Chinese (simplified - People's Republic of China).
- Chinese (traditional - Taiwan).
- French (European).
- German.
- Hebrew.
- Italian.
- Portuguese.
- Spanish (European).
- Spanish (Latin American).

DANGER

Use the supplied power cords. Ensure the facility power receptacle is the correct type, supplies the required voltage, and is properly grounded.

危险

使用所提供的电源线。确保使用正确型号的设备电源插座，提供必需的电压并且正确接地。

危險

使用隨附的電源線，確定使用正確類型的設備電源插座，提供必需的電壓，並且正確接地。

DANGER

Utiliser les câbles d'alimentation fournis. S'assurer que la prise de courant du local est du type correct, délivre la tension requise et est correctement raccordée à la terre.

GEFAHR

Die mitgelieferten Netzkabel verwenden. Sicherstellen, dass die verwendete Netzsteckdose dem vorgeschriebenen Typ entspricht, die erforderliche Spannung liefert und einwandfrei geerdet ist.

סכנה

השתמש בכבלי החשמל הנלווים. וודא כי כלי הקיבול לחשמל של המתקן הוא מהסוג הנכון, מספק את המתח הדרוש, ומוארק כהלכה.

PERICOLO

Usare il cavo di alimentazione in dotazione. Assicurarsi che la presa di corrente a disposizione sia del tipo corretto, eroghi la tensione richiesta e sia dotata di messa a terra idonea.

PERIGO

Use os cordões elétricos fornecidos. Certifique-se de que o tipo de receptor de energia da facilidade é apropriado, fornece a voltagem necessária, e está corretamente aterrado.

PELIGRO

Utilice los cables de alimentación proporcionados. Asegúrese que el receptáculo tomacorriente para la instalación sea el tipo correcto, suministre el voltaje necesario, y que esté apropiadamente puesto a tierra.

PELIGRO

Utilice los cables de alimentación proporcionados. Asegúrese que el receptáculo tomacorriente para la instalación sea del tipo correcto, suministre el voltaje necesario, y que esté apropiadamente conectado a tierra.

The following **ATTENTION** statements appear in this publication and describe practices that must be observed while installing or servicing a product. An **ATTENTION** statement provides essential information or instructions for which disregard or noncompliance may result in equipment damage or loss of data.

ATTENTION ! Activating a preferred path can result in receipt of out-of-order frames if the preferred path differs from the current path, if input and output (I/O) is active from the source port, and if congestion is present on the current path.

ATTENTION ! When configuring a PDCM array that prohibits E_Port connectivity, mistakes can render ISLs unusable and cause complex routing problems. These problems can be difficult to fault isolate and sometimes manifest incorrectly as end-device issues.

ATTENTION ! If zoning is implemented by WWN, removal and replacement of a device HBA or Fibre Channel interface (thereby changing the device WWN) disrupts zone operation and may incorrectly exclude a device from a zone.

ATTENTION ! If zoning is implemented by port number, a change to the director or switch fiber-optic cable configuration disrupts zone operation and may incorrectly include or exclude a device from a zone.

1

Introduction to McDATA Multi-Protocol Products

The enterprise-level storage area network (SAN) of today is typically complex and managed at the device layer. These problems result in SANs that use storage assets inefficiently, and are complex, error prone, expensive, and time-consuming to manage.

This chapter describes McDATA® products and SAN management applications that provide Multi-Capable Storage Network Solutions® to enable enterprises to efficiently allocate storage devices to servers on demand, while lowering the cost of storage asset ownership. These solutions:

- Consolidate information technology (IT) resources, integrate servers, interconnect remote data centers and branch offices (SAN islands), expand existing Fibre Channel SANs, and implement business continuity and disaster recovery (BC / DR) methods.
- Create a robust, reliable, scalable, and cost-effective SAN that provides enterprise-class connectivity and supports Fibre Channel Protocol (FCP) and fibre connection (FICON®) environments.
- Simplify and automate SAN management by creating an application-oriented, on-demand data network.

These solutions, coupled with McDATA's technical vision, result in greater application availability and SAN performance, thus enabling enterprises to meet their business objectives. Refer to *Chapter 4, Implementing SAN Internetworking Solutions* for detailed information.

Product Overview

McDATA provides storage network solutions that are integrated across a variety of platforms, original equipment manufacturers (OEMs), and locations. Solutions are modular and support multiple technologies (current and future), protocols, and data transmission speeds. These solutions include:

- Multi-protocol hardware.
- SAN management applications.

Multi-Protocol Hardware

McDATA provides three classes of multi-protocol hardware:

- **Intrepid®-series directors -** A director is a high port count, high-bandwidth Fibre Channel switch designed with fully-redundant, hot-swappable field replaceable units (FRUs). Directors provide superior scalability, high data security, and an availability of 99.999% (about five minutes of average down time per year). Director-class products are typically deployed at the core of large fabrics (greater than 500 ports) and are the optimum choice to support mission-critical business requirements. McDATA offers the:
  - — 64-port Intrepid 6064 Director.
  - — 140-port Intrepid 6140 Director.
  - — 256-port Intrepid 10000 Director.

  Refer to *Directors* for detailed information about each product.
- **Sphereon™-series fabric switches -** A fabric switch is a low to medium port count, high-bandwidth Fibre Channel switch designed with redundant power supplies and cooling fans. Fabric switches implement the same high-performance technology as directors, but with less redundancy, availability, and expense. Switches are cost-effective, support non-disruptive scalability and connectivity on demand, and provide an availability of 99.9% (about 8.8 hours of average down time per year). Switch-class products are typically deployed at the edge of large fabrics or provide the foundation for small (less than 200 ports) or medium fabrics (between 200 and 500 ports). McDATA offers the:

— 32-port Sphereon 3232 Fabric Switch.

— 12-port Sphereon 4300 Fabric Switch. The switch provides both switched fabric and Fibre Channel arbitrated loop (FC-AL) connectivity.

— 16-port Sphereon 4400 Fabric Switch. The switch provides both switched fabric and FC-AL connectivity.

— 24-port Sphereon 4500 Fabric Switch. The switch provides both switched fabric and FC-AL connectivity.

— 32-port Sphereon 4700 Fabric Switch. The switch provides both switched fabric and FC-AL connectivity.

Refer to *Fabric Switches* for detailed information about each product.

- **Eclipse™-series SAN routers -** A SAN router is a low port count, high-bandwidth product that unifies storage and networking architectures and provides metropolitan area network (MAN) or wide area network (WAN) extended distance access and multi-protocol access to traditional Fibre Channel SANs. McDATA offers the:

  — Four-port Eclipse 1620 SAN Router.

  — 16-port Eclipse 2640 SAN Router.

  Refer to *SAN Routers* for detailed information about each product.

McDATA products (except the Sphereon 4300 Fabric Switch) are managed and controlled through a rack-mount management server with a SAN management application installed. Multiple products and the management server communicate on a local area network (LAN) through one or more 10/100 Base-T Ethernet hubs. Each hub provides 24 Ethernet connections. Hubs are daisy-chained as required to provide additional Ethernet connections as more directors and switches are installed on a customer network.

Per customer request, directors and switches are delivered separately or installed in a McDATA FC-512 Fabricenter® equipment cabinet. The rack-mount management server is mounted at the cabinet center, and one Ethernet hub is mounted at the cabinet top. Figure 1-1 illustrates two Fabricenter equipment cabinets populated with the following:

1. Ethernet hub.
2. Sphereon 3232 Fabric Switch.

3. Intrepid 6064 Director.
4. Rack-mount management server.
5. Sphereon 4300 Fabric Switch.
6. Sphereon 4500 Fabric Switch.
7. Intrepid 6140 Director.
8. Eclipse 1620 SAN Router.
9. Eclipse 2640 SAN Router.
10. Intrepid 10000 Director.

Figure 1-1 Cabinet-Mount McDATA Products

SAN Management Applications

McDATA offers the following SAN management applications installed on the rack-mount management server:

- **SANavigator® application -** The SANavigator application (Version 4.2 or later) is an integrated software package that provides management of an enterprise-wide, heterogeneous SAN (with multiple vendor applications) from a single console.

  Element Manager applications installed on the management server are launched from the SANavigator application. These applications provide configuration, management, and status monitoring of Intrepid-series directors and Sphereon-series fabric switches. The SANavigator application also manages and monitors an entire complex fabric, including servers and storage devices from multiple OEMs that use multiple applications, protocols, and technologies. In addition, an instance of the SANvergence® Manager application can be launched from the SANavigator application.

- **Enterprise Fabric Connectivity Manager (EFCM) application -** The EFCM application (Version 8.7 or later) is an integrated software package that provides management of McDATA directors and switches from a single console.

  Element Manager applications installed on the management server are launched from the EFCM application. These applications provide configuration, management, and status monitoring of Intrepid-series directors and Sphereon-series fabric switches. In addition, an instance of the SANvergence Manager application can be launched from the EFCM application.

- **SANvergence® Manager application -** The SANvergence Manager application (Version 4.6 or later) provides management of Eclipse-series switches and advanced configuration and monitoring of servers and storage devices in a mixed-protocol (Fibre Channel and iSCSI) SAN.

  Element Manager applications installed on each Eclipse-series switch are launched from the SANvergence Manager application. These applications provide switch configuration, management, and status monitoring. The SANvergence Manager application also provides a user interface with a storage name server (SNS) database that enables auto-discovery and status monitoring of Fibre Channel and Internet protocol (IP) SAN devices.

Refer to *Chapter 2, Product Management* for information about SAN management applications and the management server. Chapter 2 also describes switch management through:

- The Internet using a product's EFCM Basic Edition interface.
- Inband (Fibre Channel) application clients.
- Simple network management protocol (SNMP) workstations.
- A command line interface (CLI) or PC-attached Telnet session.

Directors

Directors provide high-performance, dynamic connections between end devices such as servers, mass storage devices, and peripherals in a Fibre Channel switched network. Directors also support mainframe and open-systems interconnection (OSI) computing environments, and provide data transmission and flow control between device node ports (N_Ports) as dictated by the *Fibre Channel Physical and Signaling Interface* (FC-PH).

Because of high port count, non-blocking architecture, and FRU redundancy, directors offer high availability and high-performance bandwidth. Directors should be installed for:

- Backbone implementation for a large-scale enterprise SAN that requires centralized storage management, centralized backup and restore, data protection, and disaster tolerance. Refer to *General Fabric Design Considerations* for information.
- Mission-critical applications and switched data paths with no downtime tolerance.
- Performance-intense applications that require any-to-any port connectivity at a high bandwidth.

Directors also provide connectivity between servers and devices manufactured by multiple OEMs. To determine if an OEM product can communicate through connections provided by a director or if communication restrictions apply, refer to the product publications or contact McDATA.

Director Performance

Directors provide the following general performance features:

- **High bandwidth -** Ports on Intrepid-series directors provide full-duplex serial data transfer at a rate of 1.0625, 2.1250, or 10.2000 gigabits per second (Gbps).
- **Low communication overhead -** Fibre Channel protocol provides efficient use of transmission bandwidth, reduces interlocked handshakes across the communication interface, and efficiently implements low-level error recovery mechanisms. This results in little communication overhead in the protocol and a director bit error rate (BER) less than one bit error per trillion (10^{12}) bits.
- **High-availability -** To ensure an availability of 99.999%, director design provides a redundant configuration of critical components with automatic failure detection and notification. Multiple FRUs (logic cards, power supplies, and cooling fans) provide redundancy in case of failure. If an active FRU fails, the backup FRU takes over operation automatically (failover) to maintain director and Fibre Channel link operation. High availability is also provided through concurrent firmware upgrades and spare or unused Fibre Channel ports.
- **Low latency -** For 1.0625 Gbps frame traffic, the latency is less than 2.5 microseconds between transmission of a frame at a source port to receipt of the frame at the corresponding destination port (with no port contention). For 2.1250 and 10.2000 Gbps frame traffic, the latency is less than 2.0 microseconds.
- **Local control -** Actions taking place at a device N_Port seldom affect operation of other ports, therefore servers need to maintain little or no information about other connected devices in a SAN.
- **Multiple topology support -** Directors support both point-to-point and multiswitch fabric topologies and indirectly support arbitrated loop topology.
  - — Point-to-point topology provides a single direct connection between two device N_Ports. This topology supports bidirectional transmission between source and destination ports. Through dynamic switching, directors configure different point-to-point transmission paths. In all cases, connected N_Ports use 100% of the available bandwidth.

— A multiswitch fabric topology provides the ability to connect directors (and other McDATA switch elements) through expansion ports (E_Ports) and interswitch links (ISLs) to form a Fibre Channel fabric. Director elements receive data from a device and, based on the destination N_Port address, route the data through the fabric (and possibly through multiple switch elements) to the destination device.

— An arbitrated loop topology connects multiple device node loop (NL_Ports) in a loop (or hub) configuration without benefit of a multiswitch fabric. Although directors do not support direct connection of arbitrated loop devices, such devices can communicate with directors through Sphereon-series fabric switches.

- **Multiple service class support -** The Fibre Channel signaling protocol provides several classes of transmission service that support framing protocol and flow control between ports. Directors support:

  — Class 2 transmission service that provides connectionless multiplexed frame delivery service with acknowledgment. Class 2 service is best suited for mainstream computing applications.

  — Class 3 transmission service that provides connectionless, best-effort multiplexed datagram frame delivery with no acknowledgment. Class 3 service is best suited for mass storage or video applications.

  — Class F transmission service that is used by multiple directors (or fabric elements) to communicate across ISLs to configure, control, and coordinate the behavior of a multiswitch fabric.

Intrepid 6064 Director

The Intrepid 6064 Director is a second-generation, enterprise-class product that provides switched fabric connectivity for up to 64 Fibre Channel devices. The product provides high-performance scalable bandwidth, highly-available operation, redundant switched data paths, long transmission distances (up to 20 km), and high device population. Figure 1-2 illustrates the director.

Figure 1-2 Intrepid 6064 Director

The director supports McDATA's non-blocking extendable open network (EON™) architecture and concurrent firmware downloads through hot code activation (HotCAT®) technology. The director also provides a modular design that enables quick removal and replacement of FRUs, including:

- Cable management assembly and front bezel with power (green) and system error (amber) light-emitting diodes (LEDs).
- Power module assembly (with AC power switch), redundant fan modules, and redundant power supplies.
- Redundant CTP (1.0625 Gbps operation) or CTP2 (2.1250 or 10.2000 Gbps operation) logic cards.
- Redundant serial crossbar (SBAR) assembly logic cards.
- Backplane.
- A minimum of eight to a maximum of 16 Fibre Channel port cards as follows:
  - — Fiber port module (FPM) cards. Each FPM card provides four 1.0625 Gbps Fibre Channel port connections through duplex small form factor pluggable (SFP) fiber-optic transceivers.

— Universal port module (UPM) cards. Each UPM card provides four 2.1250 Gbps Fibre Channel port connections through duplex SFP fiber-optic transceivers.

— Ten-gigabit port module (XPM) cards. Each XPM card provides one 10.2000 Gbps Fibre Channel port connection through a duplex ten-gigabit small form factor pluggable (XFP) fiber-optic transceiver.

For FPM, UPM, and XPM cards, shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over singlemode fiber-optic cable. Fiber-optic cables attach to director port transceivers with duplex LC® connectors.

The power module assembly at the rear of the director also provides a 9-pin, type-D subminiature (DSUB) maintenance port for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure director network addresses.

Intrepid 6140 Director

The Intrepid 6140 Director is a third-generation, enterprise-class product that provides switched fabric connectivity for up to 140 Fibre Channel devices. The product provides high-performance scalable bandwidth, highly-available operation, redundant switched data paths, long transmission distances (up to 20 km), and high device population. Figure 1-3 illustrates the director.

The director supports McDATA's non-blocking EON architecture and concurrent firmware downloads through HotCAT technology. The director also provides a modular design that enables quick removal and replacement of FRUs, including:

- Front bezel with power (green) and system error (amber) LEDs.
- Redundant CTP (2.1250 and 10.2000 Gbps operation) logic cards.
- Redundant SBAR assembly logic cards.
- Redundant cooling fans.
- Redundant power supply and AC modules.
- Backplane.

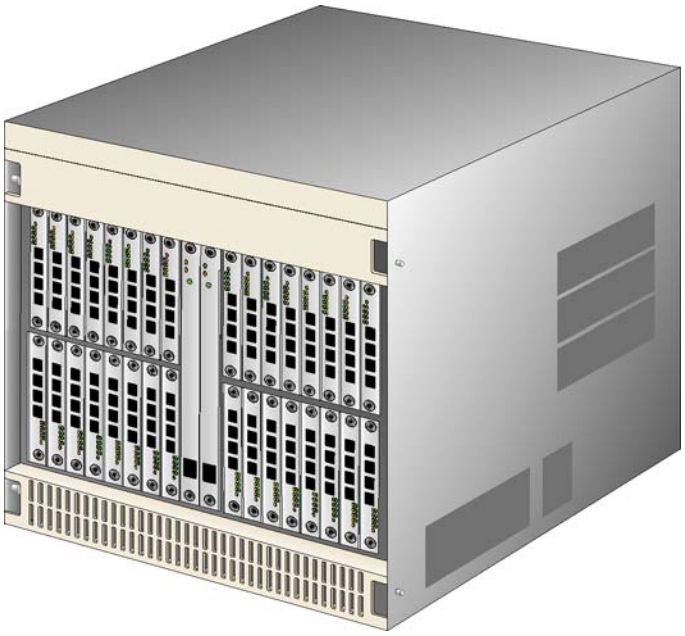

Figure 1-3 Intrepid 6140 Director

- A minimum of 16 to a maximum of 35 Fibre Channel port cards as follows:
  - — UPM cards. Each UPM card provides four 2.1250 Gbps Fibre Channel port connections through duplex SFP fiber-optic transceivers.
  - — XPM cards. Each XPM card provides one 10.2000 Gbps Fibre Channel port connection through a duplex XFP fiber-optic transceiver.

  For UPM and XPM cards, shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over singlemode fiber-optic cable. Fiber-optic cables attach to director port transceivers with duplex LC connectors.

The rear of the director provides a 9-pin, DSUB maintenance port for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure director network addresses.

Intrepid 10000 Director

The Intrepid 10000 Director is a fourth-generation, enterprise-class product that provides switched fabric connectivity for up to 256 Fibre Channel devices operating at 1.0625 or 2.1250 Gbps, or up to 64 devices operating at 10.2000 Gbps. The product provides high-performance scalable bandwidth, highly-available operation, redundant switched data paths, long transmission distances (up to 2,200 km using a pool of programmable buffer-to-buffer credits), and high device population. Figure 1-4 illustrates the director.

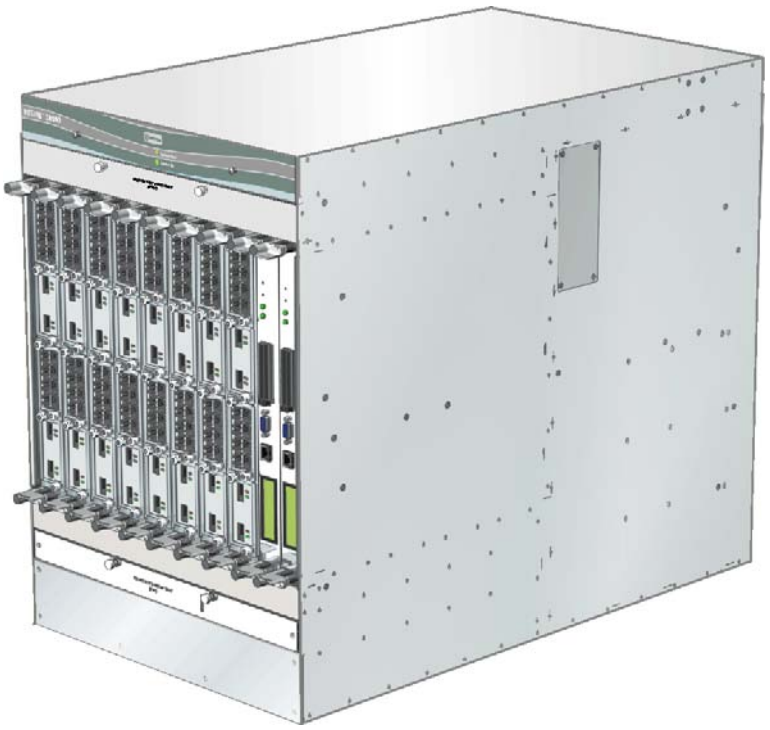

Figure 1-4 Intrepid 10000 Director

The director supports McDATA's non-blocking nScale™ architecture that allows the product to be flexibly partitioned into multiple (up to four) separate directors (FlexPar™ feature), each with its own management and Fibre Channel services subsystems. In addition, the director supports concurrent firmware downloads through HotCAT technology. FlexPar-specific software can also be independently and

concurrently upgraded. The combination of high port count and the FlexPar feature enables an enterprise to use the director at the core of both small and large SAN fabrics. For example:

- Large fabrics built around the director require fewer fabric elements (directors and switches) and ISLs. Large fabrics benefit from deterministic non-blocking performance, less ISL congestion, and better cable management. This performance is not possible from a fabric constructed with smaller port-count switches interconnected with multiple ISLs. Refer to *General Fabric Design Considerations* for information.
- Smaller fabrics or SAN islands built around the director (but separated through FlexPars) benefit from better resource utilization because port configurations are flexible to accommodate change and the hardware does not require over-provisioning for growth. The FlexPar feature enables additional fabric ports to be added to a partition on demand, without interrupting fabric traffic. Refer to *Inter-FlexPar Routing* for information.

The director provides a modular design that enables quick removal and replacement of FRUs, including:

- Front bezel with green power (**SYSTEM ON**) and amber error (**SYSTEM FAULT**) LEDs.
- Redundant CTP (1.0625, 2.1250, and 10.2000 Gbps operation) logic cards.
- Redundant and shared switching module (SWM) logic cards. The director provides a variable, full-duplex, packet-switching bandwidth that scales up to 640 Gbps (four SWMs at 160 Gbps per module).
- Redundant cooling fan trays (upper and lower, front and rear).
- Power tray with redundant, load-sharing power supplies and AC power switches.
- Upper and lower cable trays.
- A minimum of one to a maximum of eight Fibre Channel line modules (LIMs). Each LIM provides the interface to attach up to four optical paddles as follows:

— Optical paddles that operate at 1.0625 or 2.1250 Gbps provide eight Fibre Channel port connections through duplex SFP fiber-optic transceivers. A fully-populated director supports up to 256 port connections at 1.0625 or 2.1250 Gbps data rates.

— Optical paddles that operate at 10.2000 Gbps provide two Fibre Channel port connections through duplex XFP fiber-optic transceivers. A fully-populated director supports up to 64 port connections at the 10.2000 Gbps data rate.

Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over singlemode fiber-optic cable. Fiber-optic cables attach to director port transceivers with duplex LC connectors.

Fabric Switches

In similar fashion to directors, fabric switches also provide high-performance, dynamic connections between end devices in a Fibre Channel switched network. Fabric switches also support mainframe and OSI computing environments.

Through non-blocking architecture and limited FRU redundancy, fabric switches also offer high availability and high-performance bandwidth. Although switches do not offer the redundancy, availability, or port count of an enterprise-class director, they offer a much lower-cost connectivity option. Fabric switches should be installed for:

- Implementation as the principal building block of a small-scale SAN or as a consolidation point for enterprise-class SANs.
- Departmental and workgroup connectivity.
- Applications where distributed storage predominates.

Fabric switches also provide connectivity between servers and devices manufactured by multiple OEMs. To determine if an OEM product can communicate through switch connections or if communication restrictions apply, refer to the product publications or contact McDATA.

Fabric Switch Performance

Fabric switches provide an availability of 99.9% through a redundant configuration of power supplies and cooling fans. When an active FRU (power supply or fan) fails, the backup takes over operation automatically to maintain switch and Fibre Channel link operation. Availability is also provided through concurrent firmware upgrades and spare or unused Fibre Channel ports.

Excluding an availability of 99.999%, fabric switches offer the same general performance features as directors, including high bandwidth, low latency, local control, low communication overhead, multiple topology support, and multiple service class support.

Sphereon 3232 Fabric Switch

The Sphereon 3232 Fabric Switch operates at 2.1250 Gbps, provides fabric connectivity for to up to 32 Fibre Channel devices, and supports FICON, EON architecture, and HotCAT technology. Figure 1-5 illustrates the switch.

Figure 1-5 Sphereon 3232 Fabric Switch

The switch provides a modular design that enables quick removal and replacement of FRUs, including:

- Redundant power supplies and cooling fans. The switch provides two power supplies, each with an AC power receptacle and power switch. The switch also provides four cooling fans.

- Up to 32 duplex SFP fiber-optic port transceivers. Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over singlemode fiber-optic cable. Fiber-optic cables attach to switch port transceivers with duplex LC connectors.

NOTE: The Sphereon 3232 Fabric Switch can be purchased at a discount price with the Flexport Technology feature. The switch is delivered with only 16 ports enabled. When additional port capacity is required, the remaining ports are enabled (in eight-port increments) through purchase of a PFE key.

The switch front panel provides an initial machine load (**IML**) button, Ethernet LAN connector, port status LEDs, green power (**PWR**) LED, and amber system error (**ERR**) LED.

The switch rear panel provides a 9-pin DSUB maintenance port for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure switch network addresses.

Sphereon 4300 Fabric Switch

The Sphereon 4300 Fabric Switch operates at 1.0625 or 2.1250 Gbps, provides connectivity through 12 Fibre Channel ports, and supports EON architecture and HotCAT technology. Figure 1-6 illustrates the switch.

Figure 1-6 Sphereon 4300 Fabric Switch

Switch ports can be configured as:

- Fabric ports (F_Ports) to provide direct connectivity for switched fabric devices.
- Fabric loop ports (FL_Ports) to provide switched arbitrated loop connectivity and fabric attachment for FC-AL devices. The switch supports:
  - — Connectivity of public loop devices and private loop devices. Refer to *Public Versus Private Devices* for information.
  - — Configuration of public arbitrated loops and private arbitrated loops. Refer to *Public Versus Private Loops* for information.
- E_Ports to provide ISL connectivity to fabric directors and switches. E_Port connectivity is not standard and must be configured through the optional full fabric product feature enablement (PFE) key. Refer to *Full Fabric* for information.

The switch provides a modular design that enables quick removal and replacement of FRUs, including up to 12 duplex SFP fiber-optic port transceivers. Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over singlemode fiber-optic cable. Fiber-optic cables attach to switch port transceivers with duplex LC connectors.

NOTE: The Sphereon 4300 Fabric Switch can be purchased at a discount price with the Flexport Technology feature. The switch is delivered with only four ports enabled. When additional port capacity is required, the remaining ports are enabled (in four-port increments) through purchase of a PFE key.

The switch front panel provides a combined initial machine load and reset (**IML/RESET**) button, Ethernet LAN connector, port status LEDs, port speed LEDs (green for 1.0625 Gbps operation and blue for 2.1250 Gbps operation), green power (**PWR**) LED, and amber system error (**ERR**) LED.

The switch rear panel provides a 9-pin DSUB maintenance port for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure switch network addresses.

Sphereon 4400 Fabric Switch

The Sphereon 4400 Fabric Switch operates at 1.0625, 2.1250, or 4.2500 Gbps, provides connectivity through 16 Fibre Channel generic mixed ports (GX_Ports), and supports EON architecture and HotCAT technology. Figure 1-7 illustrates the switch.

Figure 1-7 Sphereon 4400 Fabric Switch

Switch ports can be configured as:

- F_Ports to provide direct connectivity for switched fabric devices.
- FL_Ports to provide switched arbitrated loop connectivity and fabric attachment for FC-AL devices. The switch supports:
  - — Connectivity of public loop devices and private loop devices. Refer to *Public Versus Private Devices* for information.
  - — Configuration of public arbitrated loops and private arbitrated loops. Refer to *Public Versus Private Loops* for information.
- E_Ports to proved ISL connectivity to fabric directors and switches.

The switch provides a modular design that enables quick removal and replacement of FRUs, including:

- Redundant power supplies and cooling fans. The switch is delivered with one external power supply assembly, and a second power supply can be installed as an option. When a second power supply is detected, the switch automatically enables high availability (HA) mode. Three internal cooling fans provide airflow.

- Up to 16 duplex SFP fiber-optic port transceivers. Tri-rate shortwave laser (1.0625, 2.1250, or 4.2500 Gbps) transceivers are available for transferring data over multimode fiber-optic cable. Tri-rate longwave laser transceivers are available for transferring data over singlemode fiber-optic cable. Fiber-optic cables attach to switch port transceivers with duplex LC connectors.

NOTE: The Sphereon 4400 Fabric Switch can be purchased at a discount price with the Flexport Technology feature. The switch is delivered with only eight ports enabled. When additional port capacity is required, the remaining ports are enabled (in four-port increments) through purchase of a PFE key.

The switch front panel provides a combined initial machine load and reset (**RESET**) button, Ethernet LAN connector, port status LEDs, green power (**PWR**) LED, and amber system error (**ERR**) LED.

The switch rear panel provides a 9-pin DSUB maintenance port for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure switch network addresses.

Sphereon 4500 Fabric Switch

The Sphereon 4500 Fabric Switch operates at 1.0625 or 2.1250 Gbps, provides connectivity through 24 Fibre Channel generic mixed ports (GX_Ports), and supports EON architecture and HotCAT technology. Figure 1-8 illustrates the switch.

Figure 1-8 Sphereon 4500 Fabric Switch

Switch ports can be configured as:

- F_Ports to provide direct connectivity for switched fabric devices.
- FL_Ports to provide switched arbitrated loop connectivity and fabric attachment for FC-AL devices. The switch supports:
  - Connectivity of public loop devices and private loop devices. Refer to *Public Versus Private Devices* for information.
  - Configuration of public arbitrated loops and private arbitrated loops. Refer to *Public Versus Private Loops* for information.
- E_Ports to proved ISL connectivity to fabric directors and switches.

The switch provides a modular design that enables quick removal and replacement of FRUs, including:

- Redundant power supplies and cooling fans. The switch provides two power supplies, each with an AC power receptacle and three cooling fans (six fans total).
- Up to 24 duplex SFP fiber-optic port transceivers. Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over singlemode fiber-optic cable. Fiber-optic cables attach to switch port transceivers with duplex LC connectors.

NOTE: The Sphereon 4500 Fabric Switch can be purchased at a discount price with the Flexport Technology feature. The switch is delivered with only eight ports enabled. When additional port capacity is required, the remaining ports are enabled (in eight-port increments) through purchase of a PFE key.

The switch front panel provides a combined initial machine load and reset (**IML/RESET**) button, Ethernet LAN connector, port status LEDs, port speed LEDs (green for 1.0625 Gbps operation and blue for 2.1250 Gbps operation), green power (**PWR**) LED, and amber system error (**ERR**) LED.

The switch rear panel provides a 9-pin DSUB maintenance port for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure switch network addresses.

Sphereon 4700 Fabric Switch

The Sphereon 4700 Fabric Switch operates at 1.0625, 2.1250, or 4.2500 Gbps, provides connectivity through 32 Fibre Channel generic mixed ports (GX_Ports), and supports FICON, EON architecture, and HotCAT technology. Figure 1-9 illustrates the switch.

Figure 1-9 Sphereon 4700 Fabric Switch

Switch ports can be configured as:

- F_Ports to provide direct connectivity for switched fabric devices.
- FL_Ports to provide switched arbitrated loop connectivity and fabric attachment for FC-AL devices. The switch supports:
  - — Connectivity of public loop devices and private loop devices. Refer to *Public Versus Private Devices* for information.
  - — Configuration of public arbitrated loops and private arbitrated loops. Refer to *Public Versus Private Loops* for information.
- E_Ports to proved ISL connectivity to fabric directors and switches.

The switch provides a modular design that enables quick removal and replacement of FRUs, including:

- Redundant power supplies and cooling fans. The switch provides two power supplies, each with an AC power receptacle and three cooling fans (six fans total).

- Up to 32 duplex SFP fiber-optic port transceivers. Tri-rate shortwave laser (1.0625, 2.1250, or 4.2500 Gbps) transceivers are available for transferring data over multimode fiber-optic cable. Tri-rate longwave laser transceivers are available for transferring data over singlemode fiber-optic cable. Fiber-optic cables attach to switch port transceivers with duplex LC connectors.

NOTE: The Sphereon 4700 Fabric Switch can be purchased at a discount price with the Flexport Technology feature. The switch is delivered with only 16 ports enabled. When additional port capacity is required, the remaining ports are enabled (in eight-port increments) through purchase of a PFE key.

The switch front panel provides a combined initial machine load and reset (**RESET**) button, Ethernet LAN connector, port status LEDs, green power (**PWR**) LED, and amber system error (**ERR**) LED.

The switch rear panel provides a 9-pin DSUB maintenance port for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure switch network addresses.

SAN Routers

The Fibre Channel protocol was designed for high-performance channel and storage applications within the limited confines of a data center. However, the protocol is not suited for long-distance applications between multiple, geographically-dispersed SANs or data centers. Conversely, transmission control protocol / Internet protocol (TCP / IP) is well suited to provide dynamic routing for complex, geographically-dispersed networks.

SAN routers provide multi-protocol solutions to this problem by unifying storage (FCP) and networking (TCP / IP) architectures. These protocols include metropolitan Fibre Channel protocol (mFCP), Internet Fibre Channel protocol (iFCP), and Internet small computer systems interface (iSCSI), provided at up to Gigabit Ethernet (GbE) bandwidth. SAN routers are low port count, high-bandwidth products that provide extended distance and multi-protocol access to Fibre Channel SANs, and should be installed to:

- **Perform SAN routing functions -** SAN routers provide FCP-protocol, router port (R_Port) connectivity between local Fibre Channel fabrics (SAN routing). A SAN routing solution provides interoperable FCP connectivity and consolidates IT resources but ensures a disruption in one fabric remains isolated and does not propagate to other fabrics. Refer to *SAN Routing* for detailed information.
- **Perform mSAN routing functions -** Multiple SAN routers interconnect with GbE bandwidth links that employ the user datagram protocol (UDP)-based mFCP protocol. These routers connect local Fibre Channel fabrics into a metropolitan storage area network (mSAN) and perform mSAN routing functions. An mSAN routing solution provides a low-latency, high-bandwidth alternative to traditional FCP connectivity. Refer to *mSAN Routing* for detailed information.
- **Implement iSAN routing and BC/DR solutions -** SAN routers provide TCP/IP-based (iFCP protocol) distance extension solutions that connect geographically-dispersed SANs into an internetworked storage area network (iSAN), perform iSAN routing, and run business continuance and disaster recovery applications over existing MAN or WAN infrastructures. Refer to *iSAN Routing* and *Implementing BC/DR Solutions* for detailed information.
- **Provide connectivity for iSCSI integration -** SAN routers provide cost-effective solutions (based on the iSCSI-protocol) to consolidate servers and storage that run a wide range of Windows-based applications. Refer to *Consolidating and Integrating iSCSI Servers and Storage* for detailed information.

SAN Router Performance

SAN routers provide the following general performance features:

- **High bandwidth -** Fibre Channel ports on the Eclipse 1620 SAN router provide full-duplex serial data transfer at a rate of 1.0625 Gbps. Fibre Channel ports on the Eclipse 2640 SAN router provide full-duplex serial data transfer at a rate of 2.1250 Gbps. Intelligent ports provide Fibre Channel data transmission or, alternately, high-speed networking (IP) bandwidth through the following:

  - **Data compression -** SAN router software identifies repetitive information in an output data stream and applies a compression algorithm to ensure the data is more compact and efficiently transmitted.
  - **FastWrite technology -** FastWrite software improves write performance over WANs by responding to initiator write commands with local transfer ready (**XFR_RDY**) commands, and buffering output data at the SAN router closest to the corresponding target device. This eliminates **XFR_RDY** command transmissions and minimizes bursty data transfer over the WAN, thus reducing round-trip delays that are characteristic of extended-distance links.
  - **Jumbo frames -** Two Ethernet frames are typically required to transmit one Fibre Channel frame consisting of 2,112 bytes. The jumbo frame feature maps one Ethernet frame to one Fibre Channel frame, thus providing more efficient data transmission.

- **High-availability -** To ensure an availability of 99.9%, multi-protocol SAN router design provides a redundant configuration of power supplies and cooling fans. High availability is also provided through concurrent firmware upgrades and spare or unused multi-protocol ports.
- **Multi-protocol support -** SAN routers support the following protocols:
  - FCP, including first, second, and third editions of the *Fibre Channel Physical and Signaling Interface* (FC-PH, PC-PH-2, and FC-PH-3), arbitrated loop (FC-AL), and R_Port. Refer to *R_Port Operation* for a discussion about R_Port operation.
  - IP networking protocols, including mFCP, iFCP, iSCSI, and Internet storage name service (iSNS). The iSNS protocol provides intelligent storage device discovery and management services comparable to those found in Fibre Channel SANs. Network protocols operate at up to full-duplex GbE bandwidth at 1,000 megabits per second (Mbps).

Eclipse 1620 SAN Router

The Eclipse 1620 SAN Router is a first-generation product that provides extended-distance, multi-protocol fabric connectivity. The primary function of the Eclipse 1620 SAN Router is to implement iFCP-based BC/DR solutions. Figure 1-10 illustrates the SAN router.

Figure 1-10 Eclipse 1620 SAN Router

SAN router ports operate as follows:

- Two user-configured FCP (**FIBRE CHANNEL 1** and **2**) ports provide 1.0625 Gbps Fibre Channel storage connectivity using SFP port connectors. FCP ports can be configured for:
  - — Auto negotiation (**FC-Auto**) operation. This is the default selection.
  - — Fabric loop port (**FL_Port**) for public loop device connectivity, fabric port (**F_Port**) for fabric device connectivity, loop port (**L_Port**) for private loop device connectivity, or router port (**R_Port**) for SAN routing operation.
- Two user-configured intelligent (**ETHERNET 3** and **4**) ports provide both FCP and IP network connectivity. Each intelligent port provides two connectors (SFP or RJ-45). Connectivity through the connector pair is mutually exclusive; only one connector can be used. Intelligent ports can be configured for:
  - — FCP storage connectivity at 1.0625 Gbps, using only the SFP port connector. The ports support FC-Auto, FL_Port, F_Port, L_Port, and R_Port operation.
  - — IP network connectivity (iFCP or iSCSI protocol) at up to full-duplex 100 Base-T Fast Ethernet (100 Mbps) port transmission speed, using only the RJ-45 port connector. Refer to *Intelligent Port Speed* for detailed information.

— IP network connectivity (iFCP or iSCSI protocol) at up to full-duplex GbE (1,000 Mbps) port transmission speed, using only the SFP port connector. Refer to *Intelligent Port Speed* for detailed information.

For SFP port connectors, shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over singlemode fiber-optic cable. Fiber-optic cables attach to port transceivers with duplex LC connectors.

SFP port transceivers are the only SAN router FRUs. The SAN router has two power supplies and eight cooling fans that are not FRUs.

The SAN router front panel provides two AC power receptacles, an Ethernet LAN connector (**MGMT**), port status LEDs, green power supply status (**PS 1** and **PS 2**) LEDs, and a green system status (**SYS**) LED. The panel also provides a 9-pin DSUB maintenance port (**MGMT**) for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure SAN router network addresses.

Eclipse 2640 SAN Router

The Eclipse 2640 SAN Router is a second-generation product that provides metropolitan or extended-distance, multi-protocol fabric connectivity. The primary function of the Eclipse 2640 SAN Router is to implement mSAN and iSAN routing solutions. Figure 1-11 illustrates the SAN router.

Figure 1-11 Eclipse 2640 SAN Router

SAN router ports operate as follows:

- Twelve user-configured FCP ports (**1** through **12**) provide 1.0625 or 2.1250 Gbps FCP storage connectivity and UDP-based network connectivity using SFP port connectors. FCP ports can be configured for:
  - — UDP-based network connectivity (mFCP protocol) at full-duplex GbE (1,000 Mbps) port transmission speed.
  - — Fibre Channel auto negotiation (**FC-Auto**) operation. This is the default selection.
  - — Fibre channel fabric loop port (**FL_Port**) for public loop device connectivity, fabric port (**F_Port**) for fabric device connectivity, loop port (**L_Port**) for private loop device connectivity, or router port (**R_Port**) for SAN routing operation.
- Four user-configurable intelligent ports (**13** through **16**) provide IP network connectivity using SFP port connectors. Intelligent ports can be configured for IP network connectivity (iFCP or iSCSI protocol) at up to full-duplex GbE (1,000 Mbps) port transmission speed. Refer to *Intelligent Port Speed* for detailed information.

The SAN router provides a modular design that enables quick removal and replacement of FRUs, including:

- Redundant power supplies and cooling fans. The SAN router provides two power supplies each with an AC power receptacle, power switch, and two cooling fans (four fans total).
- Up to 16 duplex SFP fiber-optic port transceivers. Shortwave laser transceivers are available for transferring data over multimode fiber-optic cable. Longwave laser transceivers are available for transferring data over singlemode fiber-optic cable. Fiber-optic cables attach to SAN router port transceivers with duplex LC connectors.

The SAN router front panel provides an Ethernet LAN connector (**10/100**), port status LEDs, and a green system status (**SYS**) LED. The panel also provides a 9-pin DSUB maintenance port (**CONSOLE**) for connection to a local terminal or remote terminal. Although the port is typically used by authorized maintenance personnel, operations personnel can use the port to configure SAN router network addresses.

Product Features

In addition to the characteristics and performance features described in this chapter, McDATA managed products also provide a variety of:

- Connectivity features.
- Security features.
- Serviceability features.

Connectivity Features

McDATA directors, fabric switches, SAN routers, and their associated Element Manager applications support the following connectivity features. Products or product classes that do not support a connectivity feature are noted.

- **Any-to-any connectivity -** Subject to user-defined restrictions such as zoning, directors, fabric switches, and FCP ports on SAN routers define the destination port with which a source port is allowed to communicate and provide any-to-any port connectivity. In addition, most directors and fabric switches provide connectivity for both FCP and FICON devices.

  NOTE: SAN routers do not support FICON connectivity.

- **Port blocking -** System administrators can block or unblock any director or fabric switch port through the associated Element Manager application. Blocking a port prevents an attached device from logging in to the director or switch or communicating with any attached device. A blocked port continuously transmits a Fibre Channel offline sequence (OLS).

  NOTE: SAN routers do not support port blocking.

- **Zoning -** System administrators can partition attached devices into restricted-access groups called zones. Devices in the same zone can recognize and communicate with each other through port-to-port connections. Devices in separate zones cannot recognize and communicate with each other. Directors, fabric switches, and SAN routers support port number and world-wide name (WWN) zoning.

- **State change notification -** Directors and the Sphereon 3232 Fabric Switch support a state change notification function that allows attached N_Ports to request notification when other N_Ports change operational state.

  Sphereon 4000-series fabric switches and FCP ports on SAN routers support a state change notification function that allows attached N_Ports and NL_Ports to request notification when fabric or loop-attached devices change operational state.

  Intelligent ports on SAN routers support a state change notification function using the iSNS protocol.

- **Extended distance support -** Fibre Channel ports are configured for extended distance operation (using repeaters) by changing the port buffer-to-buffer credit (BB_Credit) setting to a higher value. Refer to *Distance Extension Through BB_Credit* for detailed information. BB_Credits are configured as follows:

  — **Intrepid 6000-series Directors -** Ports configured at 1.0625 Gbps transmit data (1,800 bytes frames) up to 120 km by setting the BB_Credit value to **60**. Ports configured at 2.1250 Gbps ports transmit data (1,800 bytes frames) up to 60 km by setting the BB_Credit value to **60**. Ports configured at 10.2000 Gbps ports transmit data (2,100 bytes frames) up to 100 km by setting the BB_Credit value to **400**.

  — **Intrepid 10000 Director -** Each director LIM contains two scalable packet processors, each supporting two optical paddles (a maximum of 16 1.0625 or 2.1250 Gbps ports or four 10.2000 Gbps ports).

  After assigning BB_Credit values of **16** to short-link ports and setting the maximum BB_Credit to **1133** for a long-link port (with the remote fabric PFE key enabled), a 1.0625 Gbps port can transmit data up to 2,200 km, and a 2.1250 Gbps port can transmit data up to 1,100 km. After assigning BB_Credit values of **96** to short-link ports and setting the maximum BB_Credit to **1085** for a long-link port (with the remote fabric PFE key enabled), a 10.2000 Gbps port can transmit data up to 180 km. Refer to *Distance Extension Through BB_Credit* for configuration parameters and other detailed information.

  — **Sphereon 3232 Fabric Switch -** All 2.1250 Gbps ports transmit data (1,800 bytes frames) up to 60 km by setting the port BB_Credit value to **60**.

— **Sphereon 4300 Fabric Switch -** All switch ports are preset to a BB_Credit value of **5**. By enabling the full fabric PFE key, the switch provides a port buffer pool of 144 receive BB_Credits and the per-port value is increased to **12**, providing a data transmission distance of up to 24 km at 1.0625 Gbps and 12 km at 2.1250 Gbps. Refer to *Full Fabric* for information.

— **Sphereon 4400 Fabric Switch -** All switch ports are preset to a BB_Credit value of **6** and the switch provides a port buffer pool of 150 receive BB_Credits. Each port can be assigned between two and 120 credits, provided the total credits allocated to all ports does not exceed 150. This provides a data transmission distance of up to 240 km at 1.0625 Gbps, 120 km at 2.1250 Gbps, and 60 km at 4.2500 Gbps.

— **Sphereon 4500 Fabric Switch -** The first four ports (numbered **0** through **3**) are preset to a BB_Credit value of **12**, and the remaining ports are preset to a value of **5**. Each port can be assigned between two and 120 credits, provided the total credits allocated to all ports does not exceed 150. This provides a data transmission distance of up to 240 km at 1.0625 Gbps and 120 km at 2.1250 Gbps.

— **Sphereon 4700 Fabric Switch -** All switch ports are preset to a BB_Credit value of **12**. The switch provides two port buffer pools of 235 receive BB_Credits each (470 credits total). Ports **0** through **3**, **8** through **11**, **16** through **19**, and **24** through **27** share buffer pool zero. Ports **4** through **7**, **12** through **15**, **20** through **23**, and **28** through **31** share buffer pool one. Each port can be assigned between two and 120 credits, provided the total credits allocated to all ports in a pool does not exceed 235. This provides a data transmission distance of up to 240 km at 1.0625 Gbps, 120 km at 2.1250 Gbps, and 60 km at 4.2500 Gbps.

— **Eclipse-series SAN routers -** Intelligent ports that support IP network connectivity are not assigned BB_Credits. However, the ports provide approximately 96 megabytes (MB) of Transmission Control Protocol (TCP) packet buffering per transmission direction. TCP output buffering absorbs fabric data to ensure Fibre Channel BB_Credits are not exhausted. Up to eight MB of buffering can be allocated to any single iFCP or iSCSI session.

- **Port binding -** Directors and fabric switches support an optional feature that binds an attached Fibre Channel device to a specified product port through the device's WWN.

NOTE: SAN routers support port binding only for R_Ports.

Security Features

SAN management and Element Manager applications offer the following security features for McDATA switching products. Products or product classes that do not support a security feature are noted.

- **Password protection -** Users must provide a user name and password to log in to the management server and access all managed products. Administrators can configure user names and passwords for up to 16 users, and can authorize or prohibit specific management permissions for each user.
- **Remote user restrictions -** Remote user access to all managed products is either disabled or restricted to configured IP addresses.
- **SNMP workstation restrictions -** Remote users on SNMP workstations can only access management information base (MIB) variables managed by the product SNMP agent. SNMP workstations must belong to SNMP communities configured through the Element Manager application. If configured, the agent can send authorization failure traps when unauthorized SNMP workstations attempt to access a managed product.
- **Port blocking -** System administrators can restrict device access by blocking or unblocking any director or fabric switch port through the associated Element Manager application.

NOTE: SAN routers do not support port blocking.

- **Audit log tracking -** Configuration changes to a director or fabric switch are recorded in an audit log stored on the management server. Users can display the audit log through the Element Manager application. Log entries include the date and time of the configuration change, a description of the change, and the source of the change.

NOTE: SAN routers do not support audit log tracking.

- **Zoning -** System administrators can create zones that provide product access control to increase network security, differentiate between operating systems, and prevent data loss or corruption. Zoning can be implemented in conjunction with server-level access control and storage device access control.
- **SANtegrity® Authentication -** This feature enhances SAN security by providing password safety; challenge handshake authentication protocol (CHAP) verification for fabric elements, management servers, and devices; a product control point (PCP) user database; common transport (CT) authentication for the open-system management server (OSMS) interface; remote authentication dial-in user service (RADIUS) server support (to store and authenticate passwords and CHAP secrets); inband and out-of-band access controls lists; encrypted secure shell (SSH) protocol; and security logging.
- **SANtegrity Binding -** This feature enhances data security (in addition to SANtegrity Authentication) in large and complex SANs that are comprised of numerous fabrics and devices provided by multiple OEMs. The feature allows or prohibits director or fabric switch attachment to fabrics (fabric binding) and allows Fibre Channel device attachment to directors or fabric switches (switch binding).

NOTE: SAN routers do not support the SANtegrity Binding feature. SAN routers support port binding only for R_Ports.

Serviceability Features

McDATA directors, fabric switches, SAN routers, and the SAN management and Element Manager applications offer the following general serviceability features. Products or product classes that do not support a serviceability feature are noted.

- LEDs provide visual indicators of hardware status or malfunctions. LEDs are provided on FRUs, operator panels, front panels, and bezels.
- System alerts, event logs, audit logs, link incident logs, and hardware logs display the following director and fabric switch information at the management server or remote workstations:
  - Director status.
  - Fabric switch status.

— Ethernet link status.

— Fibre Channel link status.

In addition, threshold alerts and a threshold alert log notify users when the transmit (Tx) or receive (Rx) throughput reaches a specified value for configured ports.

- System alerts, system logs, message logs, SAN reports, and product reports display the following SAN router information at the management server or remote workstations:

  — SAN router status.

  — Port connectivity status.

  In addition, port statistics, traffic, and compression dialog boxes provide Tx or Rx status for configured ports.

- Product diagnostic software that performs power-on self-tests (POSTs), and director and fabric switch software that performs port diagnostics (internal and external loopback tests).

  NOTE: SAN routers do not support loopback testing.

- Directors and fabric switches (except the Sphereon 4300, 4400, and Sphereon 4500 Fabric Switches), can perform a diagnostic Fibre Channel (FC) wrap test. The FC wrap test applies only when a product is operated using FICON management style.

  NOTE: Sphereon 4300, 4400 and 4500 Fabric Switches do not support operation using the FICON management style.

  — Automatic notification of significant system events (to support personnel or administrators) through e-mail messages or the call-home feature.

  NOTE: SAN routers do not support the e-mail message feature. The Sphereon 4300 Fabric Switch and SAN routers do not support the call-home feature. In addition, the call-home feature may not be available if the EFC Management applications (EFCM Lite) are installed on a customer- supplied platform.

- Directors and fabric switches provide an internal modem for use by support personnel to dial in to the management server for event notification and to perform remote diagnostics.

  NOTE: SAN routers do not provide modem support.

- An RS-232 maintenance port on the director, fabric switch, or SAN router (port access is password protected) that enables installation or service personnel to change the product's IP address, subnet mask, and gateway address; or to run diagnostics and isolate system problems through a local or remote terminal.
- Redundant FRUs that are removed or replaced without disrupting product or link operation, and a modular design that enables quick removal and replacement of FRUs without the use of special tools or equipment.
- Concurrent port maintenance. FPM, UPM, and XPM cards and SFP optical transceivers are removed, added, or replaced without interrupting other ports or product operation. In addition, fiber-optic cables are attached to ports without interrupting other ports or product operation.
- Beaconing to assist service personnel in locating a specific port or product in a SAN environment. When port beaconing is enabled, the amber LED associated with the port flashes. When FRU beaconing is enabled, the amber (service required) LED on the FRU flashes. When unit beaconing is enabled, the system error LED on the product flashes. Beaconing does not affect port, FRU, or product operation.
- Status monitoring of redundant FRUs and alternate data paths to ensure continued product availability in case of failover. The SAN management application queries the status of each backup FRU daily. A backup FRU failure is indicated by an illuminated amber LED.
- Data collection through the product's Element Manager application, EFCM Basic Edition interface, or SANvergence Manager application (for SAN routers) to help isolate system problems. The data includes a memory dump file and audit, hardware, and engineering logs.

- SNMP management for directors and fabric switches using the following MIBs as defined by Internet Engineering Task Force (IETF) working documents, request for comment (RFC) memorandums, and McDATA:
  - **Fibre Channel Management Framework Integration MIB (FC-MGMT-MIB) -** This MIB (also called the Fibre Alliance MIB) defines an integrated management environment for Fibre Channel-attached devices. The MIB runs on the management server. Up to 12 authorized management workstations can be configured through the SAN management application to receive unsolicited SNMP trap messages that indicate product operational state changes and failure conditions.
  - **RFC 1213 - Management Information Base (MIB-II) for Network Management of TCP/IP-Based Internets -** This MIB defines managed objects for the Internet community and runs on each director or switch. Up to six authorized management workstations can be configured through the Element Manager application to receive unsolicited SNMP trap messages.
  - **Product-specific private enterprise MIB -** Product-specific proprietary MIBs run on each director or switch. Up to six authorized management workstations can be configured through the Element Manager application to receive unsolicited SNMP trap messages.
- SNMP management for SAN routers using the following MIBs as defined by IETF working documents, RFC memorandums, and McDATA. All listed MIBs run on each SAN router. Up to eight authorized management workstations can be configured through the Element Manager application to send SNMP trap messages that indicate product operational state changes and failure conditions. Up to four workstations can be configured to receive unsolicited SNMP trap messages.
  - **FC-MGMT-MIB -** This MIB (Fibre Alliance MIB) defines an integrated management environment for Fibre Channel-attached devices.
  - **Bridge MIB Extension Module (P-BRIDGE-MIB) -** This MIB defines objects to manage traffic-class and multicast filtering enhancements defined by IEEE 802.1D-1998.

  - **VLAN Bridge MIB Module (Q-BRIDGE-MIB) -** This MIB defines objects to manage virtual local area network (VLAN) bridging enhancements defined by IEEE 802.1Q-1998.
  - **RFC 1213 - MIB-II -** This MIB defines managed objects for the Internet community.
  - **RFC 1354 - IP Forwarding Table -** This MIB defines objects that manage IP-based Internet routing.
  - **RFC 1493 - Definitions of Managed Objects for Bridges -** This MIB defines objects that manage media access control (MAC) bridges between standard LAN segments.
  - **RFC 1757 - Remote Network Monitoring MIB -** This MIB defines objects that manage remote network monitoring devices.
  - **RFC 2851 - Textual Conventions for Internet Network Addresses -** This MIB defines text conventions that represent commonly used Internet network layer address information.
  - **Product-specific private enterprise MIB -** A variety of product-specific proprietary MIBs run on each router and contain management objects to support multi-protocol router functions.

- Advanced fabric diagnostic features that include:
  - **ISL port fencing -** Any ISL that bounces (repeatedly attempts to establish a connection) causes disruptive fabric rebuilds. ISL fencing establishes a user-defined bounce threshold that when reached, automatically blocks the disruptive E_Port.
  - **Digital SFP diagnostic support -** This feature provides access to predictive optics monitoring (POM) diagnostic data generated by newer SFP optical transceivers. The data includes temperature, transmit and receive power, and supply voltage.
  - **Embedded port log -** This log records all Fibre Channel traffic sourced from or delivered to a switch's embedded port. Log contents assist in fault diagnosis of SAN traffic problems.
  - **Embedded fabric log -** This log records events associated with the fabric controller, path selection, login server, and name server. Log contents assist in fault diagnosis of fabric problems.

2

Product Management

This chapter describes the management of McDATA multi-protocol products, including Intrepid-series directors, Sphereon-series fabric switches, and Eclipse-series SAN routers. The chapter specifically describes:

- Product management, including out-of-band (non-Fibre Channel) methods, inband (fibre connection (FICON) or Fibre Channel) methods, and a management interface summary.
- Management server support, including a description of the rack-mount management server (with specifications), associated Ethernet hub, and optional remote workstation support.
- Product firmware, including Enterprise Operating System, classic (E/OSc); Enterprise Operating System, nScale (E/OSn); and Enterprise Operating System, internetworking (E/OSi).
- Backup and restore features.
- Storage area network (SAN) management applications and associated Element Manager application graphical user interfaces (GUIs).
- The Enterprise Fabric Connectivity Manager (EFCM) Basic Edition interface.
- The command line interface (CLI).

Product Management

Out-of-band (non-Fibre Channel) management server access to all McDATA products is provided through an Ethernet local area network (LAN) connection on a director control processor (CTP) card, fabric switch front panel, or SAN router front panel. As an optional feature, inband (Fibre Channel or FICON) management access to selected McDATA products is provided through a Fibre Channel port connection.

Out-of-Band Management

The following out-of-band management access methods are provided through the management server:

- Management of directors and fabric switches through a SAN management application (SANavigator 4.2 or Enterprise Fabric Connectivity Manager (EFCM) 8.7) and associated Element Manager application. These applications are Java-based GUIs that reside on the management server under control of a Microsoft® Windows® operating system and can also be installed on remote user workstations. Refer to *SAN Management Applications* for additional information.

  Operators at remote workstations can connect to the management server through the SAN management and Element Manager applications to manage and monitor products. A maximum of 25 concurrent users can log in to the SAN management application. Refer to *Remote User Workstations* for information.

  NOTE: Product management through a SAN management and Element Manager application is not supported for the Sphereon 4300 Fabric Switch.

- Management of SAN routers through a SAN management application (SANvergence Manager 4.6) and associated Element Manager application. The SANvergence Manager application is a Java-based GUI that resides on the management server under control of a Microsoft Windows operating system. Element Manager applications installed on each router are launched from the SANvergence Manager application. Refer to *SAN Management Applications* for additional information.

- Management using simple network management protocol (SNMP). An SNMP agent is implemented through the Element Manager application that allows administrators on SNMP management workstations to access product management information using any standard network management tool. Administrators can assign Internet protocol (IP) addresses and corresponding community names as follows:

  — For directors and fabric switches, up to six workstations can be configured as SNMP trap message recipients.

  — For SAN routers, up to eight workstations can be configured as SNMP trap message originators and four workstations can be configured as SNMP trap message recipients.

  Refer to *SNMP Management Workstations* for information.

- With E/OSc, management of directors and fabric switches through the Internet using the EFCM Basic Edition interface installed on the product. This interface supports configuration, statistics monitoring, and basic operation of the product, but does not offer all the capabilities of a corresponding Element Manager application. Administrators launch the EFCM Basic Edition interface from a remote PC by entering the product's IP address as the Internet uniform resource locator (URL), then entering a user name and password at a login screen. The PC browser then becomes a management console.

  NOTE: The Intrepid 10000 Director and SAN routers do not support product management through the EFCM Basic Edition interface.

- Management of all products through a PC-based Telnet session using the CLI. Any platform that supports Telnet client software can be used.

- Management of directors and fabric switches through the EFC Management applications (EFCM Lite) shipped on a CD and installed on a customer-supplied server that meets minimum hardware requirements and uses the Microsoft Windows operating system. Contact your McDATA representative for the requirements when ordering this option.

In contrast to the applications installed on the management server, EFCM Lite does not include support for the:

— Call-home feature.

— Ability to download remote clients from the server. Install clients on remote workstations from the software distribution disk provided with this management option.

NOTE: The Sphereon 4300 Fabric Switch and SAN routers do not support product management through the EFCM Lite application.

Figure 2-1 illustrates out-of-band product management. In the figure, the managed product is an Intrepid 6064 Director. For a tabular summary of McDATA switch products and associated out-of-band management methods, refer to *Management Interface Summary*.

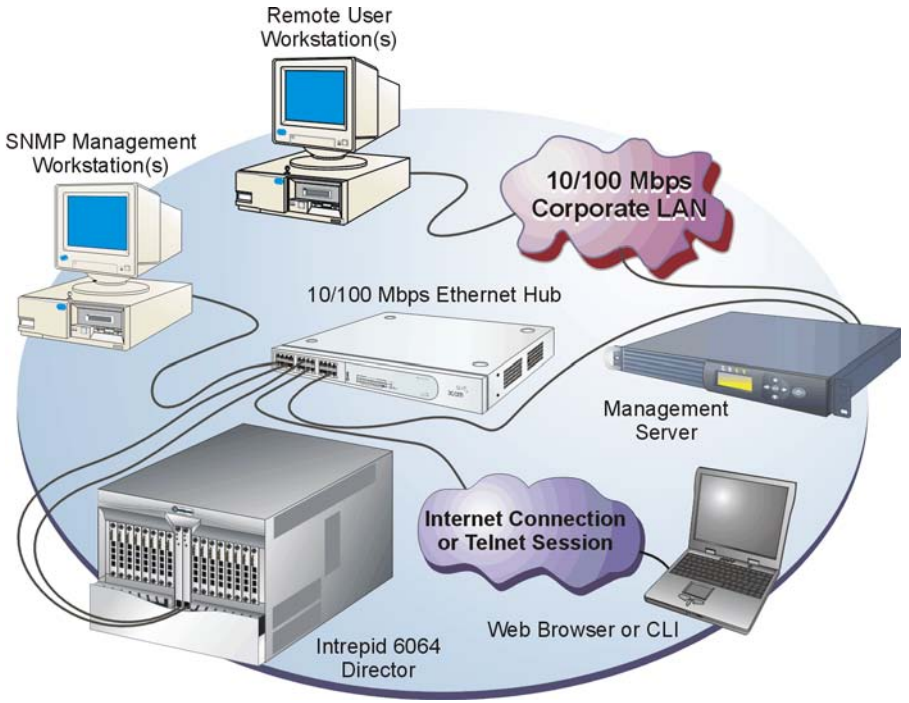

Figure 2-1 Out-of-Band Product Management

Inband Management

The following inband management access methods are provided for directors and fabric switches as options:

- Management through the product's open-system management server (OSMS) that communicates with an application client. The application resides on an open-systems interconnection (OSI) device attached to a director or switch port, and communicates using Fibre Channel common transport (FC-CT) protocol.

  Product operation, port connectivity, zoning, and fabric control are managed through a device-attached console. Refer to *OSMS* for information.

  NOTE: The Intrepid 10000 Director and SAN routers do not support out-of-band management through the OSMS.

- Management through the product's fibre connection (FICON) management server (FMS) that communicates with either the:

  — IBM® System Automation for OS/390™ (SA OS/390™) operating system resident on a System/390® (S/390) Parallel Enterprise Server™ - Generation 5 or Generation 6.

  — IBM z/OS® operating system resident on an eServer™ zSeries® 800 (z800), zSeries 900 (z900), or zSeries 990 (z990) processor.

  The server is attached to a director or switch port, and communicates through a FICON channel. Control of connectivity and statistical product monitoring are provided through a host-attached console. Refer to *FMS* for information.

  NOTE: Sphereon 4300, 4400, and 4500 Fabric Switches and SAN routers do not support out-of-band management through FMS.

Figure 2-2 illustrates inband product management. In the figure, the managed product is an Intrepid 6064 Director. For a tabular summary of McDATA switch products and associated inband management methods, refer to *Management Interface Summary*.

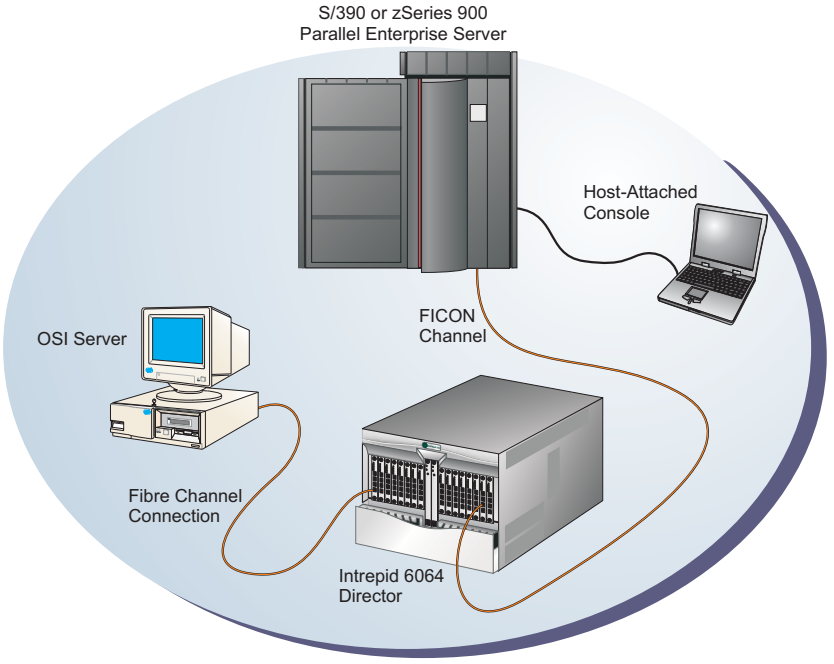

Figure 2-2 Inband Product Management

Management Interface Summary

Table 2-1 summarizes McDATA products and the out-of-band or inband management interfaces available to support the products. For each table cell, a green YES indicates the management interface supports the product, and a red NO indicates the management interface does not support the product.

Table 2-1 Out-of-Band and Inband Product Support Summary

Product	SANavigator	EFCM	SANvergence Manager	SNMP	CLI	EFCM Basic Edition	EFCM Lite	OSMS	FMS
6064 Director	YES	YES	NO	YES	YES	YES	YES	YES	YES
6140 Director	YES	YES	NO	YES	YES	YES	YES	YES	YES
10000 Director	YES	YES	NO	YES	YES	NO	YES	NO	YES
3232 Fabric Switch	YES	YES	NO	YES	YES	YES	YES	YES	YES

Table 2-1 Out-of-Band and Inband Product Support Summary (*continued*)

Product	SANavigator	EFCM	SANvergence Manager	SNMP	CLI	EFCM Basic Edition	EFCM Lite	OSMS	FMS
4300 Fabric Switch	NO	NO	NO	YES	YES	YES	NO	YES	NO
4500 Fabric Switch	YES	YES	NO	YES	YES	YES	YES	YES	NO
1620 SAN Router	NO	NO	YES	YES	YES	NO	NO	NO	NO
2640 SAN Router	NO	NO	YES	YES	YES	NO	NO	NO	NO

Management Server Support

The management server is a one rack unit (1U) high, LAN-accessed, rack-mount unit that provides a central point of control for up to 48 connected directors, fabric switches, or SAN routers. The server desktop is accessed through a LAN-attached PC and standard web browser. Figure 2-3 illustrates the server with attached liquid crystal display (LCD) panel.

Figure 2-3 Management Server

The server is rack mounted in the McDATA-supplied FC-512 Fabricenter equipment cabinet. An EFCM Basic Edition interface or management server is required to install, configure, and manage a product. Although a configured product operates normally without server intervention, an attached management server should operate at all times to monitor product operation, log events and configuration changes, and report failures.

The server is dedicated to operation of the SAN management and associated Element Manager applications. These applications provide a GUI and implement web and other server functions. Refer to *SAN Management Applications* for additional information.

NOTE: The server and SAN management application provide a GUI to monitor and manage products and are a dedicated hardware and software solution that should not be used for other tasks. McDATA tests the SAN management application installed on the server but does not compatibility test third-party software. Modifications to server hardware or installation of additional software (including patches or service packs) may interfere with normal operation.

United States English is the only language supported by the SAN management and Element Manager applications.

The server provides two auto-detecting 10/100 Mbps Ethernet LAN connectors (RJ-45 adapters). The first adapter (LAN 1) attaches (optionally) to a public customer intranet to allow access from remote user workstations. The second adapter (LAN 2) attaches to a private LAN segment containing switches or managed McDATA products.

Management Server Specifications

This section summarizes minimum and recommended hardware specifications for the rack-mount management server. Servers may ship with more enhanced hardware, such as a faster processor, additional random-access memory (RAM), or a higher-capacity hard drive.

Minimum Specifications

Minimum server specifications are:

- 1U rack-mount server running the Intel® Pentium® 4 processor with a 2 gigahertz (GHz) or greater clock speed, using the Microsoft Windows 2000 Professional (with service pack 4), Windows XP Professional (with service pack 2), or Windows Server 2003 operating system (Enterprise Edition with service pack 1) operating system.
- TightVNC™ Viewer Version 1.2.7 client-server software control package that provides remote network access (through a web browser) to the management server desktop.
- 1,024 megabyte (MB) RAM.
- 40 gigabyte (GB) internal hard drive.
- 1.44 MB 3.5-inch slim-type disk drive.

- 24X read speed slim-type compact disk-rewritable (CD-RW) and 8X read speed digital video disk (DVD) combination drive, data only.
- 56K peripheral component interconnect (PCI) internal data and fax modem, using the V .92 dial-up specification.
- 16 MB graphics card.
- Network interface card (NIC) with two 10/100 Mbps Ethernet adapters using RJ-45 connectors.

Recommended Specifications

Recommended server specifications are:

- 1U rack-mount server running the Intel Pentium 4 processor with a 3 GHz or greater clock speed, using an 800 megahertz (MHz) front side bus, using the Microsoft Windows Server 2003 operating system (Enterprise Edition with service pack 1).
- TightVNC™ Viewer Version 1.2.7 client-server software control package that provides remote network access (through a web browser) to the management server desktop.
- 2,048 MB (or greater) double data-rate synchronous dynamic random access memory (SDRAM).
- 40 GB (or greater) internal hard drive, with advanced technology attachment (ATA-100) integrated drive electronics interface operating at 7,200 rpm.
- 1.44 MB 3.5-inch slim-type disk drive.
- 48X read speed slim-type CD-RW and 32X read speed DVD combination drive, data only.
- 56K PCI internal data and fax modem, using the V .92 dial-up specification.
- Video graphics array (VGA) capable 32 MB graphics card.
- NIC with two 10/100 Mbps Ethernet adapters using RJ-45 connectors.

Ethernet Hub

The management server and managed directors, fabric switches, and SAN routers connect through a 10/100 Base-T Ethernet hub. Figure 2-4 illustrates the 24-port hub.

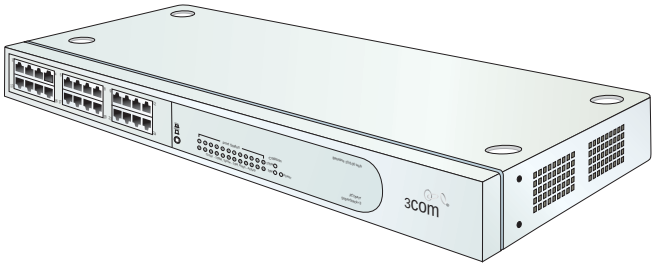

Figure 2-4 24-Port Ethernet Hub

Hubs can be daisy-chained to provide connections as additional McDATA managed products are installed on a network. Multiple hubs are daisy-chained by attaching RJ-45 Ethernet patch cables to the appropriate hub ports and configuring each hub through a medium-dependent interface (MDI) switch.

Remote User Workstations

Operators at remote workstations with client SAN management and Element Manager applications installed can connect to the management server to manage and monitor all products controlled by the server. A maximum of 25 concurrent users can log in to the SAN management application.

NOTE: The SANvergence Manager application does not support remote workstation (client) operation.

Client SAN management and Element Manager applications download and install to remote workstations (from the management server) using a standard web browser. The applications operate on platforms that meet the following minimum system requirements:

- Desktop or notebook PC with color monitor, keyboard, and mouse, using an Intel Pentium III processor with 700 MHz or greater clock speed, and using the Microsoft Windows 2000 (with service pack 4), Windows NT 4.0 (with service pack 6a), or Windows 2003 operating system.

- Unix workstation with color monitor, keyboard, and mouse, using a:
  - — Linux-based system using an Intel Pentium III processor with one GHz or greater clock speed, using the Red Hat® 7.3 or higher operating system.
  - — Hewlett-Packard® PA-RISC® processor with 400 MHz or greater clock speed, using the HP-UX® 11 or higher operating system.
  - — Sun® Microsystems UltraSPARC™ IIi or later processor, using Solaris™ Version 7.0 or higher operating system.
  - — IBM POWER3-II™ microprocessor with 333 MHz or greater clock speed, using the AIX Version 4.3.3 or higher operating system.
- At least 150 MB (Windows-based) or 350 MB (Unix-based) available on the internal hard drive.
- 512 MB or greater RAM.
- Video card supporting 256 colors at 800 x 600 pixel resolution.
- Ethernet network adapter.
- Java-enabled Internet browser, such as Microsoft Internet Explorer (Version 4.0 or later) or Netscape® Navigator (Version 4.6 or later).

Product Firmware

McDATA provides three product-embedded operating systems (firmware) that support underlying director, fabric switch, and SAN router platforms. These include:

- **E/OSc -** The Enterprise Operating System (classic) performs system configuration, management, and Fibre Channel switching functions for Intrepid 6000-series directors and Sphereon-series fabric switches.
- **E/OSn -** The Enterprise Operating System (nScale) performs system configuration, management, and Fibre Channel switching functions for the Intrepid 10000 Director.

- **E/OSi -** The Enterprise Operating System (internetworking) performs system configuration, management, and Fibre Channel and IP-based routing functions for Eclipse-series SAN routers.

Firmware Services

Director and fabric switch firmware (E/OSc and E/OSn) provides services that manage and maintain Fibre Channel connections between ports. Although product hardware transmits Fibre Channel frames between source and destination ports, the firmware maintains routing tables required by hardware to perform switching functions. SAN router firmware (E/OSi) provides services that manage and maintain both Fibre Channel and IP-based port connectivity. Product firmware also provides:

- **System management services -** This function configures, controls, and monitors product operation.
- **Application services -** This function supports all software subsystems for system initialization, logging, tracing, debugging, and communicating with RS-232 maintenance ports.
- **Operating system services -** This function includes boot and loader software, a command line monitor for engineering fault isolation, a serial maintenance port driver, and other support for the product operating system.
- **Network services -** This function provides both transmission control protocol/Internet protocol (TCP/IP) and user datagram protocol/Internet protocol (UDP/IP) transport layers to access management service subsystems from attached management clients. These clients may include (depending on the product) the out-of-band management server, EFCM Basic Edition interface, CLI, or SNMP management workstation.
- **Fibre Channel protocol services -** This function provides the Fibre Channel transport logic that allows upper layer protocols used by fabric services to communicate with devices attached to fiber-optic ports.
- **Fibre port services -** This function provides a physical driver for hardware components.

- **Fabric services -** This function supports the fabric controller (login server) and name server. For redundant directors, fabric services also implement a replication manager that synchronizes node port (N_Port) registration databases between redundant CTP cards and allows CTP failover.
- **Loop services -** This function supports FL_Port initialization for Sphereon 4000-Series Switches and implements arbitrated loop functions, such as transmission of loop initialization primitives (LIPs).
- **Port hardware services (fabric switches only) -** This function supports the application-specific integrated circuit (ASIC) embedded on the CTP card, provides frame handling for fabric switch ports, and provides the application programming interface for light-emitting diodes (LEDs), cooling fans, and power supplies.

Refer to *Appendix B, Firmware Summary* for detailed information about the differences and similarities between the operating systems. The appendix includes three tables that summarize system-related, Fibre Channel protocol-related, and management-related differences.

Backup and Restore Features

The management server provides two backup and restore features. One feature backs up (to the management server or any LAN-connected trivial file transfer protocol (TFTP) server) or restores the configuration file stored in nonvolatile random-access memory (NV-RAM) on a director CTP card, fabric switch, or SAN router. The other feature backs up (to the CD-RW drive) or restores the entire SAN management data directory. The backup and restore features operate as follows:

- **NV-RAM configuration (director or fabric switch) -** The NV-RAM configuration for any managed director or fabric switch is backed up or restored through the associated Element Manager application. Configuration data (stored in NV-RAM on each director or switch) backed up to the management server includes:
  - — Identification data, such as the product name, description, and location.
  - — Port configuration data, such as port names, port states, extended distance settings, and link incident (LIN) alerts.

  - Operating parameters, such as buffer-to-buffer credit (BB_credit), error detect timeout value (E_D_TOV), resource allocation timeout value (R_A_TOV), switch priority, switch speed (1.0625 or 2.1250 Gbps), and preferred Domain_ID.
  - Active zoning configuration.
  - SNMP configuration parameters, such as trap recipients, community names, and write authorizations.

- **NV-RAM configuration (SAN router) -** The NV-RAM configuration for any managed SAN router is backed up or restored through customer-supplied TFTP software and the associated Element Manager application.

  Configuration data (stored in NV-RAM on each SAN router) backed up to the management server is similar to configuration data backed up for directors and fabric switches. However, the management server (or any LAN-connected server on which backup is stored) requires installation of TFTP software. TFTP is a simple protocol for transferring small files, uses UDP as the transport protocol, and provides no authentication or encryption mechanisms.

- **SAN management data directory (all products) -** Critical information (for all managed products) stored in this directory is automatically backed up to a removable CD-RW when the server is rebooted or when directory contents change. The SAN management data directory includes:
  - All log files (SAN management logs and individual director or switch Element Manager logs).
  - All configuration data (product definitions, user names, passwords, user rights, nicknames, session options, SNMP trap recipients, e-mail recipients, and Ethernet event notifications).
  - Zoning library (all zone sets and zone definitions).
  - Firmware library.
  - Call-home settings (phone numbers and dialing options).
  - Configuration data for each managed product (stored on the management server and in NV-RAM on each product).

SAN Management Applications

This section describes SAN management applications that provide a GUI to monitor, manage, and control directors, fabric switches, and SAN routers. SAN management applications include SANavigator , EFCM, and SANvergence Manager. An associated Element Manager application is provided for each managed product.

NOTE: The Element Manager application for a director or fabric switch resides on the associated management server. The Element Manager application for a SAN router is a Java applet that resides on the router.

SANavigator and EFCM Applications

The management server implements a SAN management application along with director or switch-specific Element Manager applications to provide the interface for operators to control and monitor directors and fabric switches (but not SAN routers). These applications can also operate on workstations attached to the customer intranet that function as remote clients.

Application GUI

The SAN management applications provide lifecycle planning, discovery, configuration, and monitoring for an entire heterogeneous SAN. Each SAN management application is an intuitive GUI that communicates with multiple, vendor-specific applications, and provides a common tool to access the following features:

- **SAN planning -** The application provides a planning tool to develop and evaluate a SAN topology for feasibility and performance. Virtual devices and links are assembled to build a virtual SAN topology or an extension to an existing topology. The planned topology or extension is then activated to evaluate the design and identify and correct performance problems.
- **Discovery and visualization -** Through TCP/IP (out-of-band) or Fibre Channel (inband) connections, the SAN management application automatically discovers every device attached to a SAN and produces an intuitive and dynamic map of the devices and all interconnections. This map depicts device port usage, virtual and logical data paths, and allows identification of problem devices and data traffic bottlenecks.

- **Centralized configuration -** Vendor-specific device management applications can be launched from the SAN management application, including McDATA Element Manager applications. The application also provides management of director and switch zoning across multiple vendors and product models.
- **Monitoring and notification -** The application provides real-time monitoring and event notification for devices in the SAN. Informational, warning, and fatal events are recorded. The application also monitors port throughput and link performance for the entire SAN, allowing administrators to identify and solve congestion and latency issues.

The SAN management application opens automatically when the management server desktop is accessed, and the SANavigator or EFCM main window opens by default.

McDATA products, original equipment manufacturer (OEM) products, and other devices display as icons in the SANavigator main window. Only McDATA products (managed or unmanaged) display as icons in the EFCM main window. A label below each icon identifies the managed product.

For additional information about SAN management applications, refer to the *SANavigator Software Release 4.2 User Manual* (621-000013) or the *EFC Manager Software Release 8.7 User Manual* (620-000170).

Element Manager Application

An Element Manager application is provided for each managed product. The Element Manager application works in conjunction with the SAN management application and is a Java-based GUI for managing and monitoring a director or switch. The application operates locally on the management server or through a network connection from a remote PC or workstation.

NOTE: An Element Manager application is not supported for the Sphereon 4300 Fabric Switch.

To open an Element Manager application, right-click the product icon at the SAN management application's physical map, then select the *Element Manager* option from the pop-up menu. When the Element Manager application opens, the last view (tab) accessed by a user opens by default. As an example, the *Hardware View* (Figure 2-5) for the Sphereon 4500 Fabric Switch is shown.

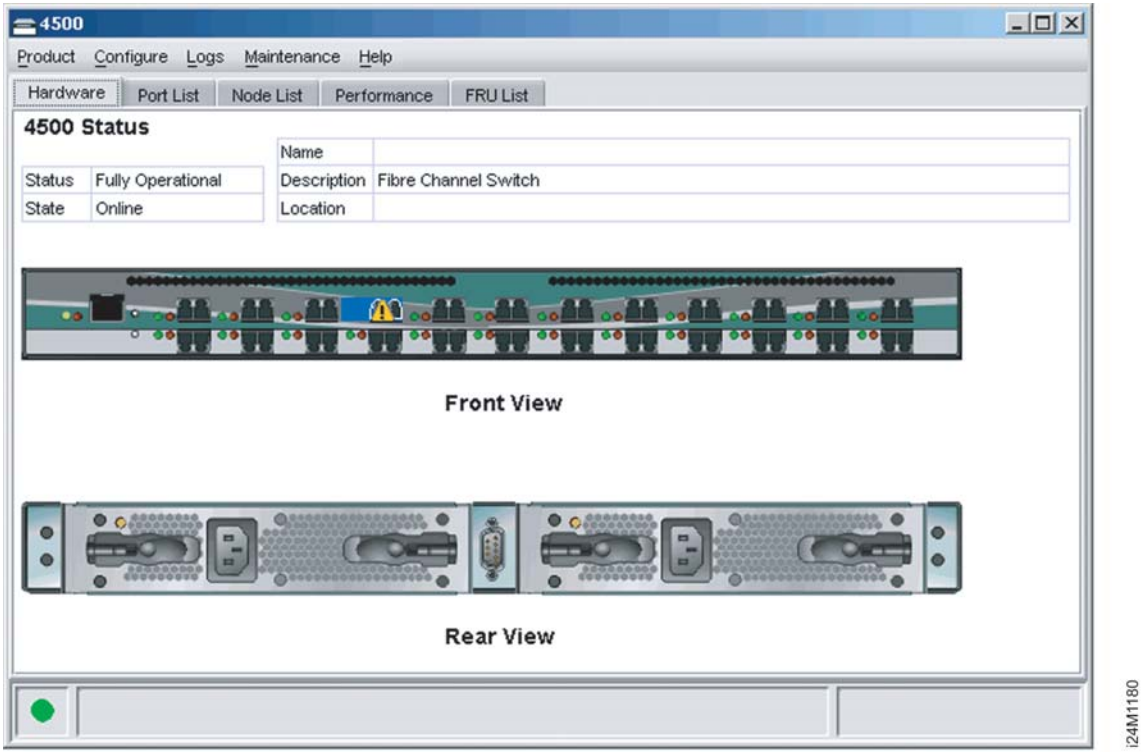

Figure 2-5 Hardware View

A status table appears at the top of the window, and a graphical representation of the hardware (front and rear) appears in the center of the window. The graphical representation of the product emulates the hardware configuration and operational status of the corresponding real product. For example, if a director or switch is fully redundant and fully populated, this configuration is reflected in the *Hardware View*. Colored symbols appear on the graphical field-replaceable units (FRUs) to represent failed or degraded status. The light-emitting diodes (LEDs) also highlight to emulate real LED operation.

A menu bar at the top of the *Hardware View* provides *Product*, *Configure*, *Logs*, *Maintenance*, and *Help* options (with associated pop-up menus) that allow users to perform Element Manager tasks.

A status bar at the bottom left corner of the view window displays colored icons (green circle, yellow triangle, red and yellow diamond, or grey square) that indicate the status of the selected managed product. Messages display as required to the right of the icons.

SANvergence Manager Application

This section describes the SANvergence Manager and Element Manager applications that provide a GUI to monitor and manage SAN routers, attached Fibre Channel elements, and metropolitan storage area network (mSAN) connectivity. An Element Manager application is provided for Eclipse 1620 and 2640 SAN Routers.

Application GUI

The SANvergence Manager application is an intuitive GUI that communicates with an attached metropolitan simple name server (mSNS). Through the mSNS, the application auto-discovers all SAN Routers, directors, and fabric switches in the mSAN; monitors product operational status, and reports problems in an event log. The application is opened from the management server desktop. When the application starts, the main window opens (Figure 2-6).

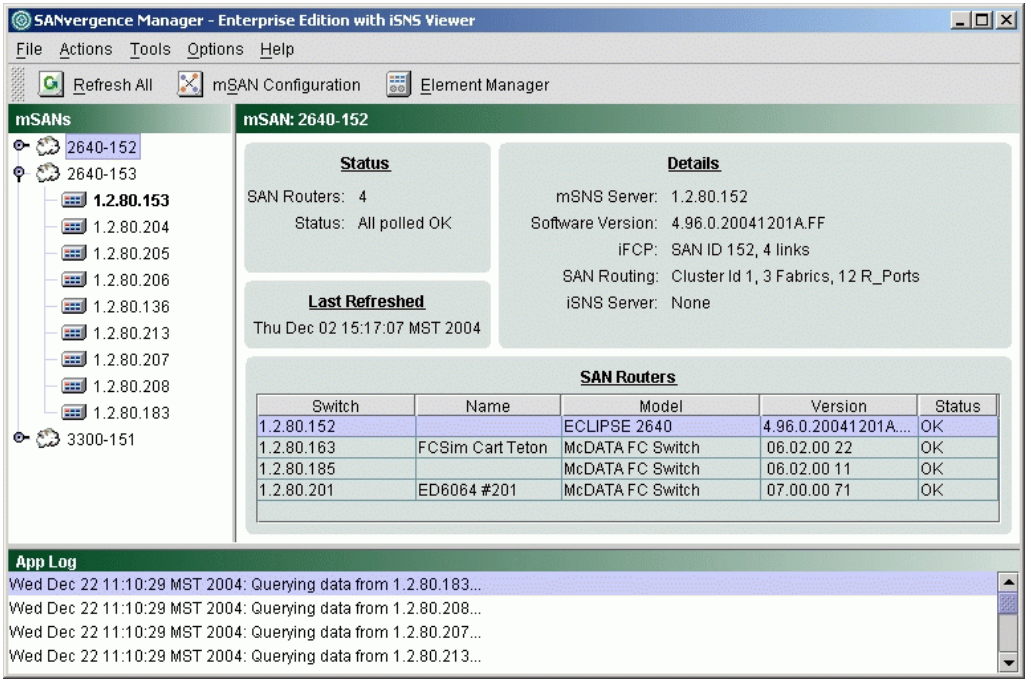

Figure 2-6 Main Window (SANvergence Manager)

For additional information about the application, refer to the *McDATA SANvergence Manager User Manual* (620-000189).

Element Manager Application

An Element Manager application is provided for each managed SAN router. The application works in conjunction with the SANvergence Manager application and is a router-resident, Java-based applet for managing and monitoring the product.

To open an Element Manager application, select (highlight) the product at the *mSANs* panel and click the *Element Manager* button on the toolbar. When the Element manager application opens, the device window opens (Figure 2-7). An Eclipse 2640 SAN Router is shown.

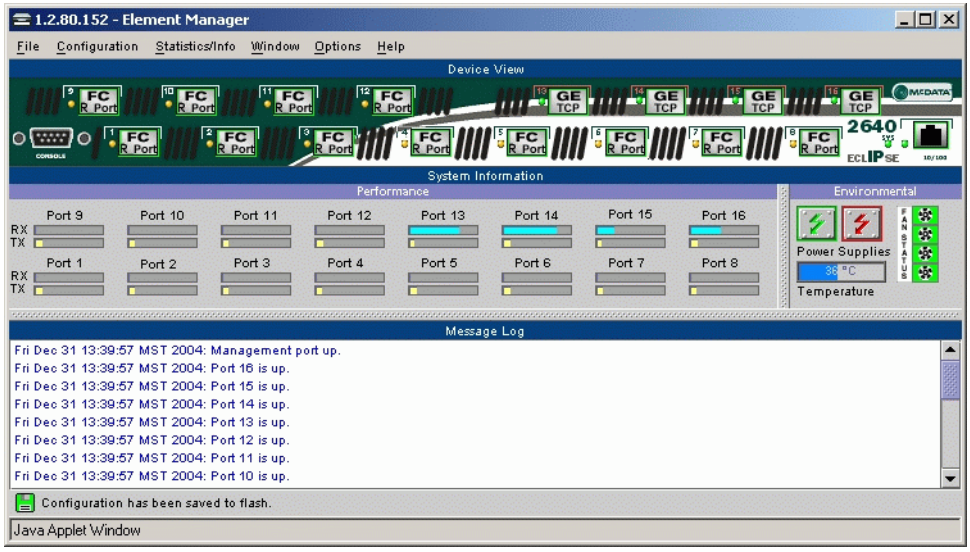

Figure 2-7 Device Window (Element Manager)

The graphical representation of the product emulates the hardware configuration and operational status of the corresponding real product. For example, if all router ports are connected and functional, this configuration is reflected in the device window. Mouse selections (right or left click) open dialog boxes or menus that display FRU properties or allow users to perform operations and maintenance tasks.

Colored symbols appear on the graphical FRUs to represent failed or degraded status. The LEDs also highlight to emulate real LED operation.

A menu bar at the top of the device window provides *File*, *Configuration*, *Statistics/Info*, *Window*, *Options*, and *Help* selections (with associated pop-up menus) that allow users to perform Element Manager tasks.

EFCM Basic Edition Interface

With E/OSc firmware installed, administrators or operators with a browser-capable PC and Internet connection can monitor and manage a product through the EFCM Basic Edition interface.

The interface is opened from a standard web browser running Netscape Navigator 4.6 or higher or Microsoft Internet Explorer 4.0 or higher. At the browser, enter the IP address of the product as the Internet uniform resource locator (URL). When prompted at a login screen, enter a user name and password. When the interface opens, the default display is the *Hardware View* (Figure 2-8). The *Hardware View* for the Sphereon 4500 Fabric Switch is shown as an example.

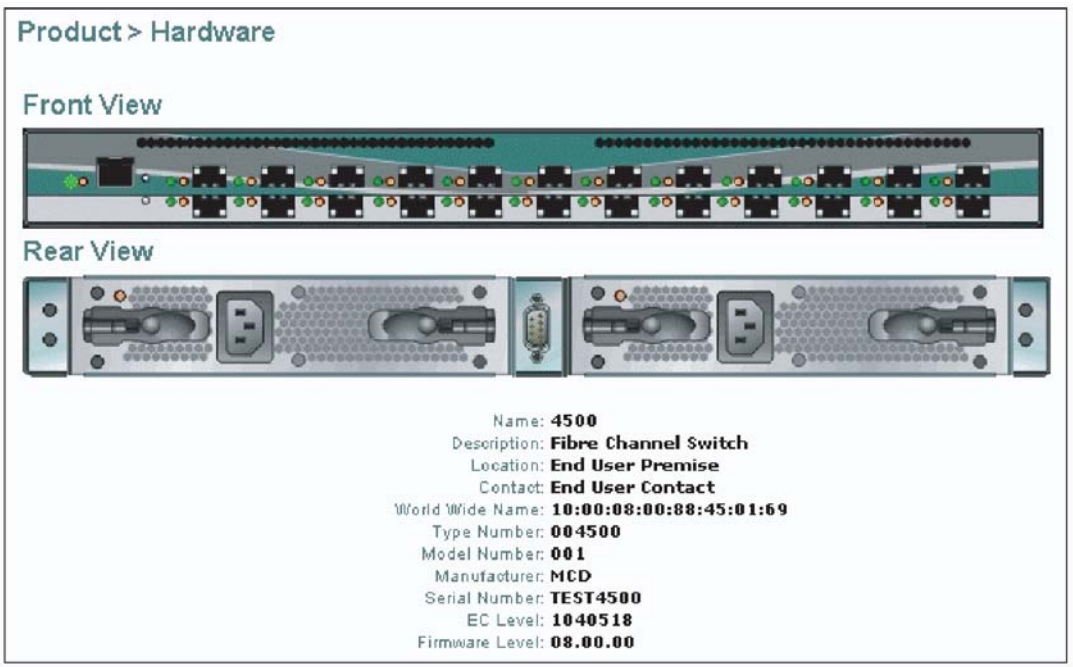

Figure 2-8 Hardware View (EFCM Basic Edition Interface)

The interface provides a GUI similar to the Element Manager application and supports product configuration, statistics monitoring, and basic operation. The EFCM Basic Edition interface manages only a single product (but has hyperlink access to other switches in a fabric). For additional information about the application, refer to the *McDATA EFCM Basic Edition User Manual* (620-000240).

Command Line Interface

The CLI provides a director, fabric switch, and SAN router management alternative to traditional SAN management GUIs. The interface allows users to access application functions by entering commands through a PC-attached Telnet session. Any platform that supports Telnet client software can be used.

The primary purpose of the CLI is to automate management of several directors or switches using scripts. Although the CLI is designed for use in a host-based scripting environment, basic commands can be entered directly at a disk operating system (DOS) window command prompt. The CLI is not an interactive interface; no checking is done for pre-existing conditions, and a user prompt does not display to guide users through tasks.

For additional information, refer to the following publications:

- *McDATA E/OSc Command Line Interface User Manual* (620-000134). This publication describes CLI support for Intrepid 6000-series directors and Sphereon 4000-series fabric switches.
- *McDATA E/OSn Command Line Interface User Manual* (620-000211). This publication describes CLI support for the Intrepid 10000 Director.
- *McDATA E/OSi Command Line Interface User Manual* (620-000207). This publication describes CLI support for Eclipse-series SAN routers.

3

Planning Considerations for Fibre Channel Topologies

A storage area network (SAN) is typically defined as a network of shared storage resources that can be allocated throughout a heterogeneous environment. This chapter describes planning considerations for incorporating McDATA products into Fibre Channel SAN topologies. The chapter specifically describes:

- Fibre Channel topologies (arbitrated loop and multiswitch fabric).
- Characteristics of arbitrated loop operation.
- Planning for private arbitrated loop connectivity.
- Planning for fabric-attached arbitrated loop connectivity.
- Fabric topologies (mesh, core-to-edge, and SAN islands).
- Planning for multiswitch fabric support.
- General fabric design considerations.
- Large fabric design considerations.
- Mixed fabric design considerations.
- Fibre connection (FICON) cascading.

Fibre Channel Topologies

Intrepid-series directors and Sphereon-series fabric switches support device connectivity through multiswitch fabric topologies. Sphereon 4000-series fabric switches also support connectivity through an arbitrated loop topology. A combination of these topologies (hybrid topology) is also supported. Topologies are described as follows:

- **Arbitrated loop -** This topology uses a Sphereon 4000-series fabric switch to connect multiple device node loop ports (NL_Ports) in a Fibre Channel arbitrated loop (FC-AL) or hub configuration without benefit of a multiswitch fabric. Both switches support a switched mode topology that provides a single, logical connection between two device NL_Ports. The switches dynamically configure different logical transmission paths, and in all cases, connected NL_Ports have access to 100% of the available bandwidth.

  Loop devices communicate with switches through a fabric loop port (FL_Port). If peripheral loop devices are expected to communicate with fabric-attached devices, consider installation of a Sphereon 4000-series fabric switch to form a fabric-loop hybrid topology. For information, refer to *Planning Considerations for Fibre Channel Topologies*, *Planning for Private Arbitrated Loop Connectivity*, and *Planning for Fabric-Attached Loop Connectivity*.

- **Multiswitch fabric -** This topology provides the ability to connect directors and fabric switches through expansion ports (E_Ports) and interswitch links (ISLs) to form a Fibre Channel fabric. Director or fabric switch elements receive data from a device and, based on the destination N_Port address, route the data through the fabric (and possibly through multiple switch elements) to the destination device. For additional information, refer to *Planning for Multiswitch Fabric Support* and *General Fabric Design Considerations*.

Characteristics of Arbitrated Loop Operation

When implementing Fibre Channel arbitrated loop topology, consideration must be given to switch operating mode, device connectivity, and loop configuration. This section describes the characteristics of arbitrated loop operation, including:

- Switch operation in shared mode or switched mode.
- Connectivity of public loop devices and private loop devices.
- Configuration of public arbitrated loops and private arbitrated loops.

This section focuses on loop operation for Sphereon 4000-series fabric switches that operate at 1.0625, 2.1250, or 4.2500 gigabits per second (Gbps) and support FC-AL operation using FL_Ports and public and private device connectivity.

Shared Mode Versus Switched Mode

Legacy arbitrated loop switches (such as the McDATA ES-1000 Switch) are configured to operate in user-selectable shared or switched mode.

NOTE: Sphereon 4000-series fabric switches do not support shared mode operation.

Shared Mode Operation

When set to shared mode, a switch acts as a hub implementing standard Fibre Channel arbitrated loop topology (although the loop has the physical appearance of a star configuration) and distributes the frame routing function through each loop port. When a loop circuit is initialized and established, arbitration protocol ensures only one device attached to a hub port (H_Port) owns the loop at a time. The port establishes communication with another device attached to an H_Port and allows the devices to transmit or receive frames. During frame transmission between these devices, the full bandwidth of the switch is used and no other H_Ports or devices are available for connection. When frame transmission completes, the loop circuit closes and other devices are able to contend for operation (using standard loop arbitration). Shared mode operation and its simplified logical equivalent are illustrated in Figure 3-1.

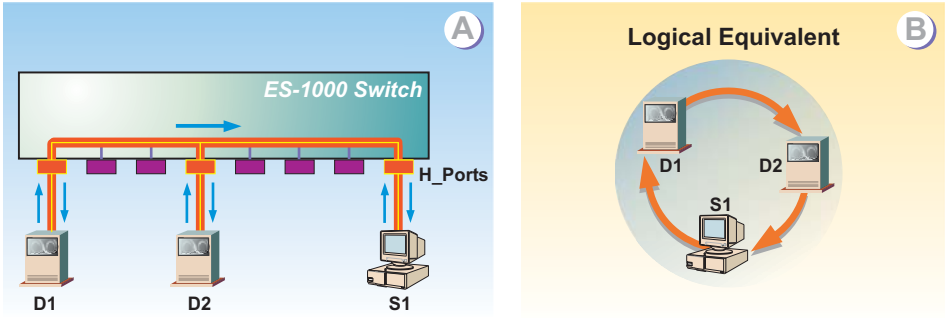

Figure 3-1 Shared Mode Operation and Logical Equivalent

Part (A) of Figure 3-1 shows device D_1 connected to server S_1 through a pair of H_Ports. Although the remaining switch H_Ports (six ports) and device D_2 are unavailable for connection, frame traffic between device D_1 and server S_1 travels through a loop that consists of all eight H_Ports, device D_1, device D_2, and server S_1. Each H_Port not participating in the communication pair and the NL_Port on device D_2 provide a repeater function that allows frames to pass around the loop at the full switch bandwidth.

Part (B) of Figure 3-1 shows the logical equivalent of this arbitrated loop. When frame transmission between device D_1 and server S_1 completes, the loop circuit closes and other ports attached to initiating devices arbitrate for loop access. When operating in shared mode, the switch is a serially reusable resource that provides service access to all ports on the loop. Access is granted by successful arbitration. When arbitration is won by a device, the loop is busy and other arbitrating devices must wait for loop access.

Switched Mode Operation

When set to switched mode, a switch bypasses full loop arbitration and enables frame transmission through multiple point-to-point connected pairs. Switched mode operation and its simplified logical equivalent are illustrated for a Sphereon 4500 Fabric Switch in Figure 3-2.

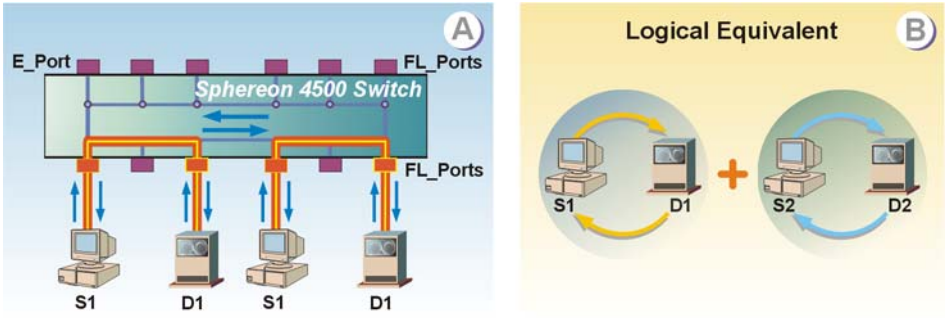

Figure 3-2 Switched Mode Operation and Logical Equivalent

Part (A) of Figure 3-2 shows server S_1 connected to device D_1 through a switched pair of FL_Ports communicating at 1.0625 Gbps. Server S_2 is connected to device D_2 through a second switched pair of ports, also communicating at 1.0625 Gbps. Because of opportunistic bandwidth sharing, the two switched pairs effectively increase the switch bandwidth to 2.1250 Gbps. The remaining ports are available for switched connection to each other but cannot communicate with servers S_1 and S_2 or devices D_1 and D_2. Part (B) of Figure 3-2 shows the logical equivalent of this arbitrated loop.

Switched mode operation provides the ability to design a complex and high-performance SAN for the department or workgroup. Consider the following when planning such a SAN:

- Connect loop switch ports to multiple unmanaged hubs to provide additional FC-AL device connectivity in the form of looplets. Cascade the unmanaged hubs if more hubs are necessary for the configuration.
- Attach devices that frequently communicate with each other to the same looplet to take advantage of opportunistic bandwidth sharing (communication predominately stays within the loop). Switched connections allow connectivity as necessary to devices attached to other looplets.
- Each looplet acts as a normal FC-AL loop. Spread multiple servers and high bandwidth storage devices across several looplets to avoid performance problems associated with a single looplet.
- Consider data traffic capacity of the department or workgroup (normal and peak load) as part of the switch planning and installation process. Capacity planning:
  - — Ensures loop traffic is distributed and balanced across servers and storage devices.
  - — Identifies traffic bottlenecks and provides for alternate connectivity solutions if required.
  - — Assists in calculating scalability to satisfy nondisruptive growth requirements or eventual connection to a Fibre Channel switched fabric.

  Capacity planning is a dynamic activity that must be performed when new devices, applications, or users are added to the department or workgroup loop configuration.

Public Versus Private Devices

Sphereon 4000-series fabric switches support connection of public and private arbitrated loop devices as follows:

- **Public device -** A loop device that can transmit a fabric login (FLOGI) command to the switch, receive acknowledgement from the switch's login server, register with the switch's name server, and communicate with fabric-attached devices is a public device.

  The switches provide loop connectivity up to the Fibre Channel architectural limit of 127 devices per Fibre Channel port when configured as an FL_Port. Each FL_Port is assigned one arbitrated loop physical address (AL_PA), leaving 126 AL_PAs per port available for device connections. Up to 32 public devices can be connected to each of the FL_Ports. As shown in Figure 3-3, server S_2 is a public loop device connected to a Sphereon 4500 Fabric Switch and can communicate with fabric-attached device D_1.

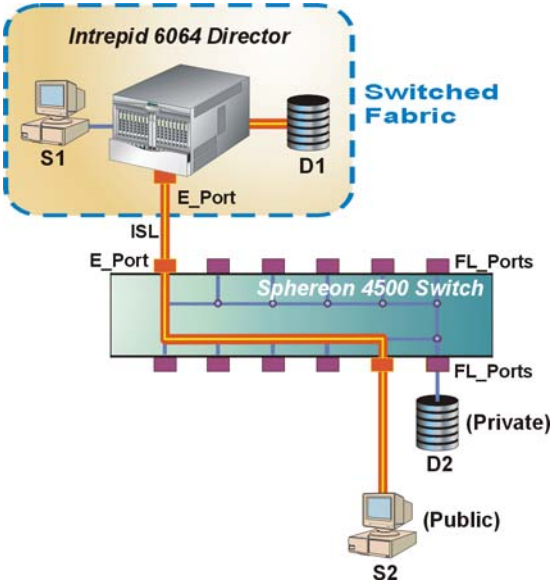

Figure 3-3 Public Device Connectivity

Public devices support normal fabric operational requirements, such as fabric busy and reject conditions, frame multiplexing, and frame delivery order.

- **Private device -** A loop device that cannot transmit an FLOGI command to the switch nor communicate with fabric-attached devices is a private device. As shown in Figure 3-4, device D_2 is a private loop device connected to a Sphereon 4500 Fabric Switch and cannot communicate with any fabric-attached device. However, device D_2 can communicate with switch-attached server S_2 (using private addressing mode).

  Public and private devices are partitioned into two separate address spaces defined in the Fibre Channel address. Private address spaces are isolated from a switched fabric. The switch does not support any other form of Fibre Channel address conversion (spoofing) that would allow private device-to-fabric device communication.

  NOTE: A private device can connect to the switch (loop) while a public device is connected and using an E_Port to communicate with a switched fabric.

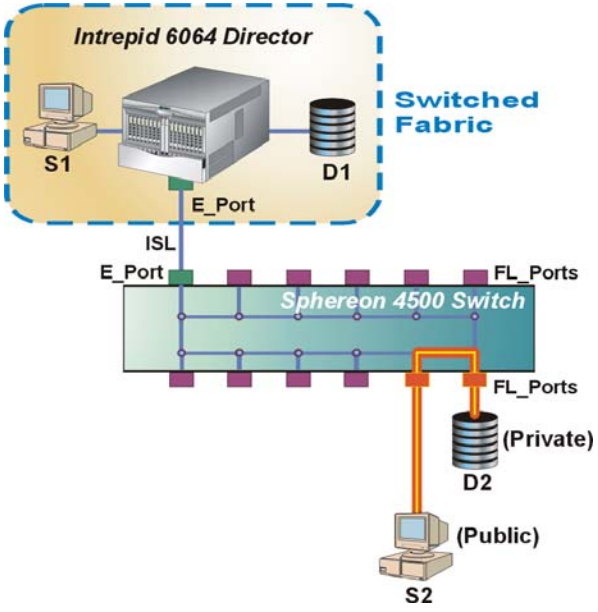

Figure 3-4 Private Device Connectivity

Private devices only communicate with other devices on the same arbitrated loop, and interconnected public and private devices can communicate with each other. Such intermixed devices establish operating parameters and loop topology configuration through a port login (PLOGI) command exchange, rather than through the switch's name server.

Be aware that public device-to-private device communication may cause problems. For example, it is often critical to separate servers and storage devices with different operating systems because accidental transfer of information from one to another can delete or corrupt data. Plan to implement security provisions for the switch, such as partitioning attached devices into restricted-access groups (zoning), providing server-level access control (persistent binding), or providing storage-level access control. Refer to *Security Provisions* for additional information.

Public Versus Private Loops

Sphereon 4000-series fabric switches support operation of public and private loops as follows:

- **Public loop -** A public loop is connected to a switched fabric through any active FL_Port. All devices attached to the loop can communicate with each other, and public devices attached to the loop can communicate with fabric-attached devices connected:
  - — Directly to another switch port configured as a fabric port (F_Port).
  - — Another fabric director or switch connected to the fabric switch through any active E_Port.

  Public loop connectivity for a Sphereon 4500 Fabric Switch is illustrated in Figure 3-5.

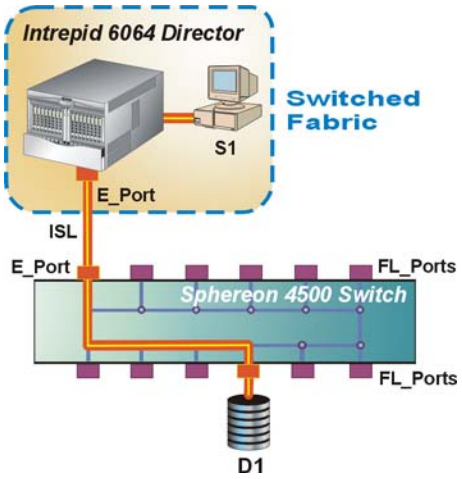

Figure 3-5 Public Loop Connectivity

- **Private loop -** A private loop is not connected to a switched fabric and the switch's embedded FL_Port is inactive. All devices attached to the loop can only communicate with each other. Private loop connectivity for a Sphereon 4500 Fabric Switch is illustrated in Figure 3-6.

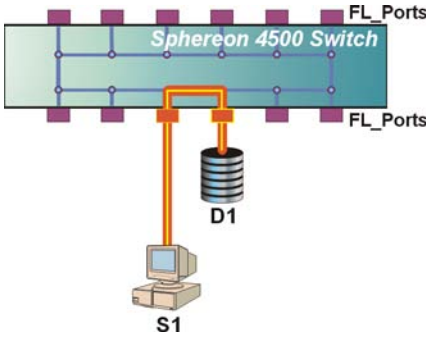

Figure 3-6 Private Loop Connectivity

FL_Port Connectivity

Sphereon 4000-series fabric switches provide loop connectivity through GX ports that are active as FL_Ports. The ports provide port addressing, physical connectivity, and Fibre Channel frame routing. Each FL_Port (and the embedded FL_Port) has a 24-bit address identifier. The address identifier is expressed in hexadecimal format as ***DD AA PP***, where:

- ***DD*** is the domain identifier. This identifier is assigned one of 31 values (**60** through **7F**).
- ***AA*** is the area identifier. Each physical FL_Port (12 or 24 ports) is assigned one of up to 24 values (**04** through **1B**).
- ***PP*** is the port identifier. Each device (node) attached to an FL_Port is assigned one of 126 AL_PA values (**01** through **EF**). The embedded FL_Port is assigned an AL_PA of **00**.

Planning for Private Arbitrated Loop Connectivity

Private arbitrated loop topology supports the clustering of isolated servers and storage subsystems into workgroup or departmental SANs. This topology is well-suited to small and mid-sized configurations where modest connectivity levels and high data transmission speeds are required. The topology also supports low-cost switching and connectivity in environments where the per-port cost of a switched fabric director is prohibitive. Private arbitrated loop topology:

- Supports the connection of up to 126 NL devices per loop plus the switch's embedded FL_Port (127 connections).
- Reduces connection costs by distributing the routing function through each loop port (loop functionality is a small addition to normal Fibre Channel port functionality).
- Provides a fully-blocking architecture that allows a single connection between any pair of loop ports. Connections between a third loop port and busy ports are blocked until communication between the first connection pair ends.

Planning for Fabric-Attached Loop Connectivity

Public arbitrated loop topology supports the connection of workgroup or departmental FC-AL devices to a switched fabric through any Sphereon 4000-series fabric switch port active as an E_Port. This topology is well-suited to:

- Providing connectivity between a workgroup or departmental SAN and a switched fabric, thus implementing connectivity of FC-AL devices to fabric devices at the core of the enterprise.
- Consolidating low-cost Windows NT or Unix server connections and providing access to fabric-attached storage devices.
- Consolidating FC-AL tape device connections and providing access to fabric-attached servers.

NOTE: For the Sphereon 4300 Fabric Switch, E_Port connectivity is not standard and must be configured through an optional product feature enablement (PFE) key

Connecting FC-AL Devices to a Switched Fabric

Sphereon 4000-series fabric switches provide dynamic connectivity between FC-AL devices and directors or switches participating in a switched fabric. This function allows multiple low-cost or low-bandwidth departmental or workgroup devices to communicate with fabric-attached devices through a high-bandwidth link and provides connectivity as required to an enterprise SAN environment. This approach provides:

- Cost-effective FC-AL device connectivity to a switched fabric. A loop switch provides fabric connectivity without incurring true switched fabric costs.
- Improved access and sharing of data and computing resources throughout an organization by connecting isolated departmental or workgroup devices to the core data center. Fabric-to-loop connectivity ensures edge servers have access to enterprise storage, and edge peripherals have access to enterprise computing resources.

- Improved resource manageability. Distributed resources are consolidated and managed through Fibre Channel connectivity instead of physical relocation. One management server manages the operation and connectivity of multiple fabric directors, fabric-attached devices, arbitrated loop switches, and FC-AL devices.
- Improved security of business applications and data. Fabric directors and a loop switch allow fabric-attached and FC-AL devices to be partitioned into restricted-access zones to limit unauthorized access. Refer to *Zoning* for information.

When using a fabric switch to provide loop-to-switched fabric connectivity and incorporate FC-AL devices into the enterprise SAN environment, attach the device to any switch port. The port senses the FC-AL device and initializes itself as an FL_Port. The loop device can communicate with fabric devices attached directly to the switch (connected through F_Ports), or with fabric devices connected to another switch and communicating through an E_Port (ISL).

Server Consolidation

Providing fabric connectivity for multiple low-bandwidth servers (Windows NT or Unix-based) by attaching them individually to an expensive Fibre Channel director is not a cost-effective solution. A practical solution is to consolidate servers on an inexpensive loop switch, then connect the switch to a director E_Port. Figure 3-7 illustrates the consolidation of ten servers on a Sphereon 4500 Fabric Switch (using two unmanaged hubs) through one E_Port connection to a fabric director.

Each server has a ten MBps bandwidth, therefore the sum of the bandwidths of all consolidated servers equals the E_Port bandwidth of 1.0625 Gbps. Connecting another server to the switch would exceed the E_Port capability and adversely impact director-to-switch link performance. Other devices (such as tape drives) should not be connected to a switch used for server consolidation.

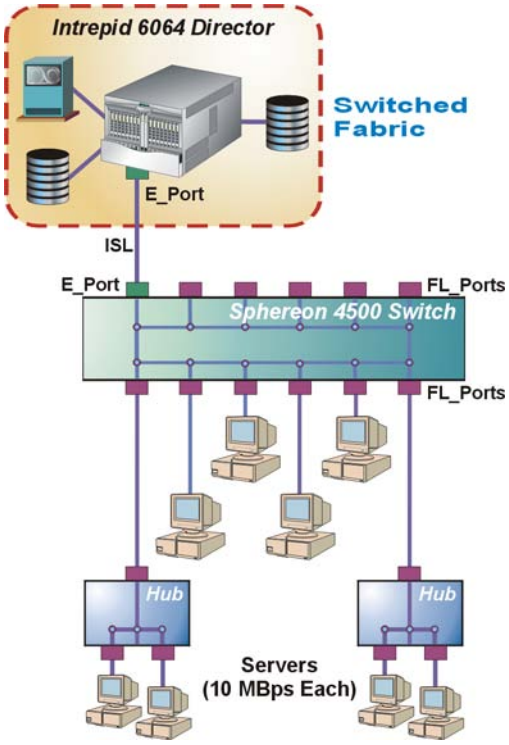

Figure 3-7 Server Consolidation

Tape Device Consolidation

Providing fabric connectivity for multiple FC-AL tape drives by attaching them individually to a Fibre Channel director is likewise not a cost-effective solution. A practical solution is to consolidate the tape drives on an inexpensive loop switch, then connect the switch to a director E_Port.

Figure 3-8 illustrates the consolidation of three tape drives through one E_Port connection to a fabric director. Each tape drive has a 30 MBps bandwidth, therefore the sum of the bandwidths of all consolidated servers is slightly less (90 MBps) than the E_Port bandwidth of 1.0625 Gbps. Connecting another FC-AL tape drive to the switch would exceed the E_Port capability and adversely impact director-to-switch link performance. Other devices (such as servers) should not be connected to a switch used for tape drive consolidation.

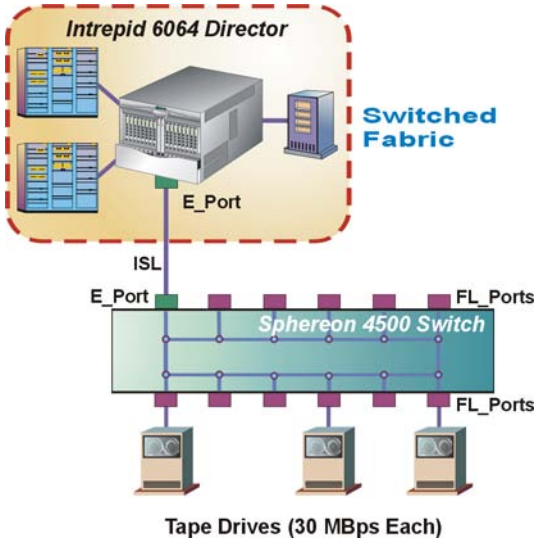

Figure 3-8 Tape Drive Consolidation

Fabric Topologies

Several topologies exist from which to build a Fibre Channel fabric infrastructure. This section describes the most effective fabric topologies and provides guidance on when to deploy each topology. The topologies are effective for a wide variety of applications, are extensively tested by McDATA, and are deployed in several customer environments. Fabric topologies described in this section include:

- Mesh.
- Core-to-edge.
- Fabric (SAN) islands.

Mesh Fabric

There are two types of mesh fabrics: full mesh and modified (or partial) mesh. In a full-mesh topology, every director or switch is directly connected to all directors and switches in the fabric. The maximum hop count between fabric-attached devices is one hop. Figure 3-9 illustrates the topology.

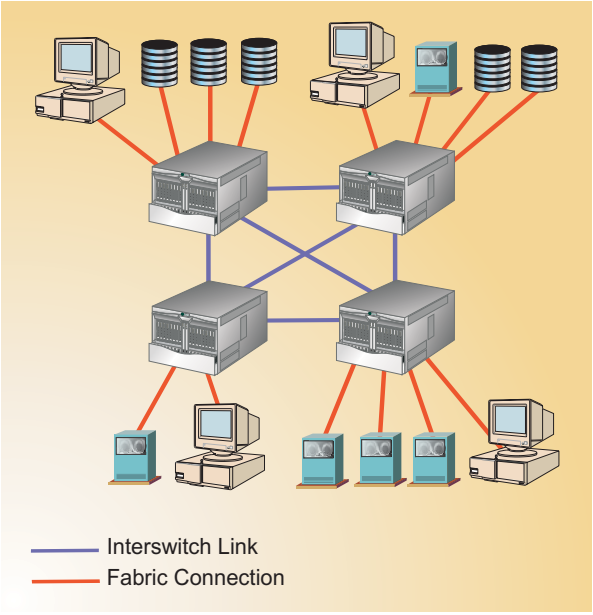

Figure 3-9 Full Mesh Fabric

Full-mesh fabrics provide increased resiliency over cascaded or ring fabrics and are well suited for applications that require any-to-any connectivity. If a single ISL fails, traffic is automatically routed through an alternate path.

Mesh fabrics also form effective backbones to which other SAN islands can be connected. Traffic patterns through the fabric should be evenly distributed and overall bandwidth consumption low.

When using low port-count fabric elements, mesh fabrics are best used when the fabric is not expected to grow beyond four or five switches. The cost of ISLs becomes prohibitive for larger mesh fabrics. In addition, full-mesh fabrics do not scale easily because the addition of a switch requires that at least one additional ISL be added from every existing switch in the fabric. If less than four fabric elements are used in a full-mesh fabric:

- A two-switch full mesh fabric is identical to a two-switch cascaded fabric.
- A three-switch full mesh fabric is identical to a three-switch ring fabric.

A modified or partial-mesh fabric is similar to a full-mesh fabric, but each switch does not have to be directly connected to every other switch in the fabric. The fabric is still resilient to failure but does not carry a cost premium for unused or redundant ISLs. In addition, partial-mesh fabrics scale easier than full-mesh fabrics. Partial-mesh fabrics are useful when designing a SAN backbone for which traffic patterns between SAN islands connected to the backbone are well known. If heavy traffic is expected between a pair of switches, the switches are connected through at least one ISL; if minimal traffic is expected, the switches are not connected.

In general, mesh fabrics can be difficult to scale without downtime. The addition of switches or directors usually involves disconnecting fabric devices and may involve disconnecting in-place ISLs. As a result, full or partial-mesh fabrics are recommended for networks that change infrequently or have well-established traffic patterns.

Core-to-Edge Fabric

A core-to-edge fabric consists of one or more directors or switches acting as core elements that are dedicated to connecting other directors and switches (edge elements) in the fabric. Core directors act as high-bandwidth routers with connectivity to edge fabric elements. Figure 3-10 illustrates the topology with two core directors and fourteen edge directors and switches (2-by-14 topology).

Subject to large fabric design constraints, core-to-edge fabrics are easy to scale through the addition of core elements. The topology offers any-to-any device connectivity and evenly distributes traffic bandwidth throughout the fabric. The topology provides the most flexible architecture to address fabric performance, traffic locality, data integrity, connectivity, and scalability requirements.

The simplest core-to edge fabric has two or more core switching elements that may or may not be connected (simple or complex). In a simple core topology as shown in Figure 3-10, core switches are not connected. In a complex core topology, core switches are connected. The figure also illustrates a topology where the core is a full-mesh fabric.

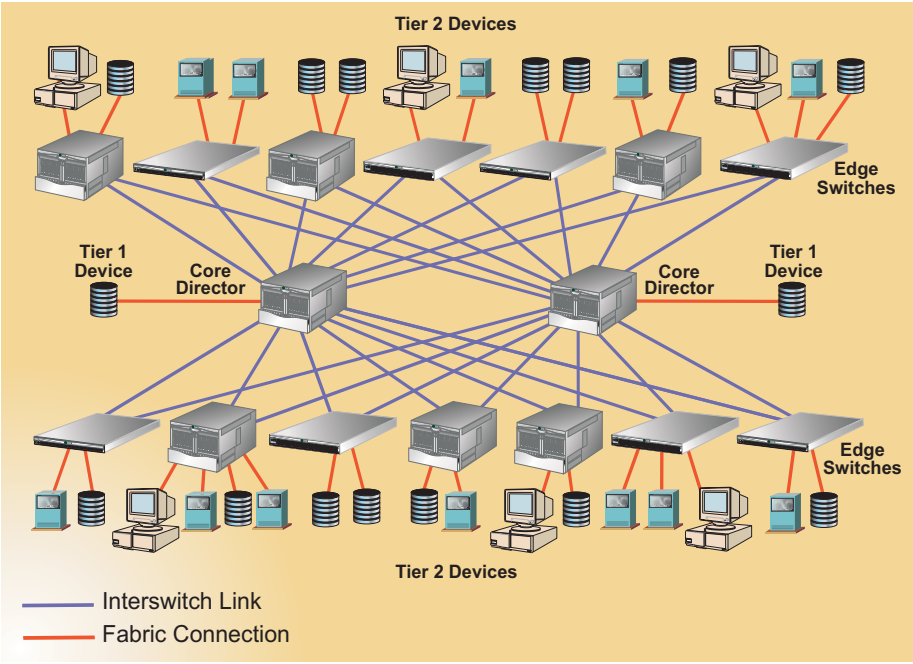

Figure 3-10 2-by-14 Core-to-Edge Fabric

Each edge switch connects (through at least one ISL) to each core switch but not to other edge switches. There are typically more device connections to an edge switch than ISL connections; therefore edge switches act as consolidation points for servers and storage devices. The ratio of ISLs to device connections for each switch is a function of device performance. For additional information, refer to *ISL Oversubscription*.

Fibre channel devices (servers and storage devices) connect to core or edge fabric elements in tiers. These tiers are defined as follows:

- **Tier 1 -** A Tier 1 device connects directly to a core director or switch. Tier 1 devices are typically high-use or high-I/O devices that consume substantial bandwidth and should not be connected through an ISL. In addition, fibre connection (FICON) devices cannot communicate through E_Ports (ISLs) and must use Tier 1 connectivity. For additional information, refer to *FCP and FICON in a Single Fabric*.

- **Tier 2 -** A Tier 2 device connects to an edge switch and Fibre Channel traffic from the device must traverse only one ISL (hop) to reach a device attached to a core director or switch.
- **Tier 3 -** A Tier 3 device connects to an edge switch and Fibre Channel traffic from the device can traverse two ISLs (hops) to reach a device attached to a core director or switch.

SAN Islands

A SAN island is an isolated or geographically diverse Fibre Channel fabric. These fabrics may also comprise different topologies (mesh or core-to-edge), but may require connectivity for shared data access, resource consolidation, data backup, remote mirroring, or disaster recovery.

When connecting multiple fabrics, data traffic patterns and fabric performance requirements must be well known. Fabric island connectivity must adhere to topology limits, including maximum number of fabric elements and ISL hop count. It is also essential to maintain data locality within fabric islands as much as possible and to closely monitor bandwidth usage between the fabric islands. Refer to *SAN Island Consolidation* for additional information.

Planning for Multiswitch Fabric Support

A Fibre Channel topology that consists of one or more interconnected director or switch elements is called a fabric. The product operational software provides the ability to interconnect directors and switches (through E_Port connections) to form a multiswitch fabric. Support of multiswitch fabric operation is a major feature of a director or fabric switch. Consider installation of multiple directors or switches to form a high-availability fabric topology that supports multiple, full-bandwidth data transmission paths between servers and devices. Figure 3-11 illustrates a simple multiswitch fabric. In the figure, the three fabric elements are Intrepid 6064 Directors.

Fabric elements cooperate to receive data from the N_Port of an attached device, route the data through the proper director or switch fabric ports (F_Ports), and deliver the data to the N_Port of a destination device. The data transmission path through the fabric is typically determined by the fabric elements and is transparent to the user. Subject to zoning restrictions, devices attached to any of the interconnected directors or switches can communicate with each other through the fabric.

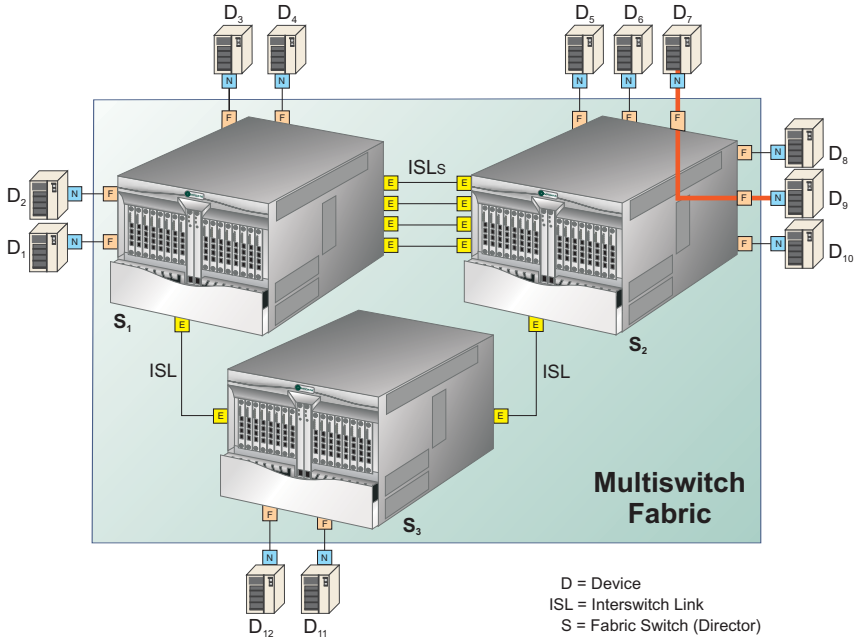

Figure 3-11 Example Multiswitch Fabric

A multiswitch fabric is typically complex and provides the facilities to maintain routing to all device N_Ports attached to the fabric, handle flow control, and satisfy the requirements of the classes of Fibre Channel service that are supported.

Fabric Topology Limits

Operation of multiple directors or switches in a fabric topology is subject to the following topology limits. Consider the impact of these limits when planning the fabric.

- **Fabric Elements -** Each fabric element is defined by a unique domain identification (Domain_ID) that ranges between **1** and **31**. A Domain_ID of **0** is invalid. Therefore, the theoretical limit of interconnected directors or switches supported in a single fabric is **31**. For additional information, refer to *General Fabric Design Considerations*.

- **Heterogeneous fabric -** Vendor interoperability in the fabric environment is supported; therefore, fabric elements can include directors, fabric switches, and open-fabric compliant products supplied by original equipment manufacturers (OEMs). To determine if interoperability is supported for a product or if communication restrictions apply, refer to the supporting publications for the product or contact McDATA.
- **Number of ISLs -** The Intrepid 6064 Director supports 48 ISLs. The Intrepid 6140 Director supports 140 ISLs. The Intrepid 10000 Director supports seven ISLs per optical paddle pair. Sphereon-class switches support a maximum ISL count equal to half the number of Fibre Channel ports available on the product. For redundancy, at least two ISLs should connect any pair of director-class fabric elements.
- **Hop count -** The Fibre Channel theoretical limit of ISL connections traversed (hop count) in a single path through the fabric is seven. The tested and verified hop count limit is three.

NOTE: The hop count is equal to the number of ISL connections traversed in a single path, not the total number of ISL connections between devices. As shown in Figure 3-11, the number of ISL connections between switch **S1** and **S2** is four, while the number of hops is one.

Factors to Consider When Implementing a Fabric Topology

Director and switch-based fabrics offer scalable, high-performance, and high-availability connectivity solutions for the enterprise. To enable a multiswitch fabric, all fabric elements must be defined to the SAN management application and must be physically cabled to form the requisite ISL connections. In addition, it is recommended that each director or switch in the fabric be assigned a unique preferred Domain_ID. When planning to implement a fabric topology, consider the following connectivity and cabling best practices:

- **Physical characteristics and performance objectives -** Most enterprises have unique configurations determined by the characteristics of end devices, fabric elements, cost, and the installation's performance objectives (such as high data transfer rate or high availability). These factors, along with nondisruptive growth and service requirements, must be evaluated when planning an initial fabric. For additional information, refer to *General Fabric Design Considerations*.

- **Distance requirements -** The distance between elements in a fabric affects the type of optical port transceiver and cabling required. In addition, variables such as the number of connections, grade of fiber-optic cable, device restrictions, application restrictions, buffer-to-buffer credit limits, and performance requirements can affect distance requirements. Consider the following:

  — If the distance between two fabric elements is less than 300 meters (at 1.0625 Gbps), 150 meters (at 2.1250 Gbps), 75 meters (at 4.2500 Gbps), or 33 meters (at 10.2000 Gbps) any port transceiver (shortwave or longwave laser) and any fiber-optic cable type (50-micron multimode, 62.5-micron multimode, or 9-micron singlemode) can be used to create an ISL. Cost or port availability may be the determining factor.

  — If the distance between two fabric elements is between 300 and 500 meters (at 1.0625 Gbps), 150 and 300 meters (at 2.1250 Gbps), 75 and 150 meters (at 4.2500 Gbps), or 33 and 82 meters (at 10.2000 Gbps) any port transceiver (shortwave or longwave laser) and 50-micron multimode or 9-micron singlemode fiber-optic cable can be used to create an ISL.

  — If the distance between two fabric elements exceeds 500 meters (at 1.0625 Gbps), 300 meters (at 2.1250 Gbps), 150 meters (at 4.2500 Gbps), or 82 meters (at 10.2000 Gbps) only longwave laser port transceivers and 9-micron singlemode fiber-optic cable can be used to create an ISL.

  — Distance limitations can be increased by using multiple fabric elements. Each director or switch retransmits received signals, thus performing a repeater and multiplexer function. However, be aware that each connection introduces a nominal signal loss of at least one dB through the ISL. If dB losses introduced through multiple connections exceed the link budget of the entire ISL, link errors occur. Refer to *Data Transmission Distance* for additional information about link budgets and distance limitations.

    Distance limitations can also be increased by using a variety of local area network (LAN), metropolitan area network (MAN) or wide area network (WAN) extension technologies. For additional information, refer to *SAN Extension Transport Technologies*.

- **Bandwidth -** ISL connections can be used to increase the total bandwidth available for data transfer between two directors or switches in a fabric. Increasing the number of ISLs between elements increases the corresponding total ISL bandwidth but decreases the number of port connections available to devices.
- **Load balancing -** Planning consideration must be given to the amount of data traffic expected through the fabric or through a fabric element. Because the fabric automatically determines and uses the least cost (shortest) data transfer path between source and destination ports, some ISL connections may provide insufficient bandwidth while the bandwidth of other connections is unused.

  To optimize bandwidth use and automatically provide dynamic load balancing across multiple ISLs, consider purchasing and enabling the OpenTrunking feature key. For information about the feature and managing multiple ISLs, refer to *OpenTrunking* and *General Fabric Design Considerations*.
- **Preferred path -** Preferred path is an option that allows a user to configure an ISL data path between multiple fabric elements (directors and fabric switches) by configuring the source and exit ports of the origination fabric element and the Domain_ID of the destination fabric element. Each participating director or switch must be configured as part of a desired path. For information about the feature, refer to *Preferred Path*.

ATTENTION ! Activating a preferred path can result in receipt of out-of-order frames if the preferred path differs from the current path, if input and output (I/O) is active from the source port, and if congestion is present on the current path.

In general, Fibre Channel frames are routed through fabric paths that implement the minimum possible hop count. For example, in Figure 3-11, all traffic between devices connected to director $\mathbf{S_1}$ and director $\mathbf{S_2}$ communicate directly through ISLs that connect the directors (one hop). No traffic is routed through director $\mathbf{S_3}$ (two hops). If heavy traffic between the devices is expected, multiple ISL connections should be configured to create multiple minimum-hop paths. With multiple paths, the directors balance the load by assigning traffic from different ports to different minimum-hop paths (ISLs).

When balancing a load across multiple ISLs, a director or switch attempts to avoid assigning multiple ports attached to a device to the same ISL. This minimizes the probability that failure of a single ISL will affect all paths to the device. However, because port assignments are made incrementally as devices log into the fabric and ISLs become available, optimal results are not guaranteed.

Special consideration must also be given to applications with high data transfer rates or devices that participate in frequent or critical data transfer operations. For example, in Figure 3-11, suppose device $\mathbf{D_7}$ is a server and device $\mathbf{D_9}$ is a storage unit and both devices participate in a critical nightly backup operation. It is recommended that such a connection be routed directly through director $\mathbf{S_2}$ (rather than the entire fabric) through zoned port connections, WWN-bound port connections, or a preferred path. For additional information, refer to *Device Locality*.

- **Zoning -** For multiswitch fabrics, zoning is configured on a fabric-wide basis. Changes to the zoning configuration apply to all directors and switches in the fabric. To ensure the zoning configuration is maintained, certain rules are enforced when two or more elements are connected through ISLs to form a fabric or when two or more fabrics are joined. For additional information, refer to *Configuring Zones*.

After directors and fabric switches are defined and cabled, they automatically join to form a single fabric through a user-transparent process. However, the user should be aware of the following fabric concepts, configuration characteristics, and operational characteristics:

- **Principal switch selection -** Setting this value determines the principal switch for the multiswitch fabric. Select either *Principal* (highest priority), *Default*, or *Never Principal* (lowest priority) from the *Switch Priority* drop-down list.

  If all fabric elements are set to *Principal* or *Default*, the director or fabric switch with the highest priority and the lowest WWN becomes the principal switch. Following are some examples of principal switch selection when fabric elements have these settings.

  — If you have three fabric elements and set all to *Default*, the director or switch with the lowest WWN becomes the principal switch.

— If you have three fabric elements and set two to *Principal* and one to *Default*, the element with the *Principal* setting that has the lowest WWN becomes the principal switch.

— If you have three fabric elements and set two to *Default* and one to *Never Principal*, the element with the *Default* setting and the lowest WWN becomes the principal switch.

Note that at least one director or switch in a multiswitch fabric needs to be set as *Principal* or *Default*. If all the fabric elements are set to *Never Principal*, all ISLs will segment. If all but one element is set to *Never Principal* and the element that was *Principal* goes offline, then all of the other ISLs will segment.

NOTE: It is recommended to configure the switch priority as *Default*.

In the audit log, note the *Principal* setting maps to a number code of **1**, *Default* maps to a number code of **254**, and *Never Principal* maps to a number code of **255**. Number codes **2** through **253** are not used.

- **Fabric WWN assignment -** The Fabric Manager application identifies fabrics using a fabric WWN. The fabric WWN is the same as the WWN of the fabric's principal switch. If a new principal switch is selected because of a change to the fabric topology, the fabric WWN changes to the WWN of the newly selected principal switch.
- **Domain_ID assignment -** Each director or switch in a multiswitch fabric is identified by a unique Domain_ID that ranges between **1** and **31**. A Domain_ID of **0** is invalid. Numerical Domain_IDs specified by a user are converted to hexadecimal format and used in 24-bit Fibre Channel addresses that uniquely identify source and destination ports in a fabric.

  Each fabric element is configured through the Element Manager application with a preferred Domain_ID. When a director or switch powers on and comes online, it requests a Domain_ID from the fabric's principal switch (indicating its preferred value as part of the request). If the requested Domain_ID is not allocated to the fabric, the Domain_ID is assigned to the requesting director or switch. If the requested Domain_ID is already allocated, an unused Domain_ID is assigned.

If two operational fabrics join, they determine if any Domain_ID conflicts exist between the fabrics. If one or more conflicts exist, the interconnecting ISL E_Ports segment to prevent the fabrics from joining. To prevent this problem, it is recommended that all directors and switches be assigned a unique preferred Domain_ID. This is important if zoning is implemented through port number (and by default Domain_ID) rather that WWN.

When assigning preferred Domain_IDs in an open fabric with directors and switches supplied by multiple OEMs, be aware of the following:

— For Intrepid 6000-series directors and Sphereon-series fabric switches, the firmware adds a base offset of **96** (hexadecimal **60**) to the numerically-assigned preferred Domain_ID. Therefore, if a user assigns a director or switch a numerical preferred Domain_ID of **1**, the firmware assigns a hexadecimal Domain_ID of **61**.

— For the Intrepid 10000 Director and Eclipse-series SAN routers, the firmware does not add a base offset to the numerically-assigned preferred Domain_ID.

— For non-McDATA directors and switches, the product firmware may not add a base offset to the numerical preferred Domain_ID or may add a different hexadecimal base offset (not **60**).

As a consequence of this variable base offset and hexadecimal conversion, Domain_ID conflicts may exist in an open fabric, even if each participating director and switch is assigned a unique numerical Domain_ID. To determine the method of preferred Domain_ID assignment for a product, refer to the supporting OEM publications for the product or contact McDATA.

NOTE: Do not assign Domain_ID **30** or Domain_ID **31** to a fabric element. In a routed SAN, these proxy Domain_IDs are assigned to routing domains.

- **Path selection -** Directors and fabric switches are not manually configured with data transmission paths to each other. Participating fabric elements automatically exchange information to determine the fabric topology and resulting minimum-hop data transfer paths through the fabric. These paths route Fibre

Channel frames between devices attached to the fabric and enable operation of the fabric services firmware on each director or switch.

Paths are determined when the fabric topology is determined and remain static as long as the fabric does not change. If the fabric topology changes (elements are added or removed or ISLs are added or removed), directors and switches detect the change and define new data transfer paths as required. The algorithm that determines data transfer paths is distributive and does not rely on the principal switch to operate. Each director or switch calculates its own optimal paths in relation to other fabric elements.

Only minimum-hop data transfer paths route frames between devices. If an ISL in a minimum-hop path fails, directors and switches calculate a new least-cost path (which may include more hops) and route Fibre Channel frames over that new path. Conversely, if the failed ISL is restored, directors and switches detect the original minimum-hop path and route Fibre Channel frames over that path.

When multiple minimum-hop paths (ISLs) between fabric elements are detected, firmware balances the data transfer load and assigns ISL as follows:

— The director or switch assigns an equal number of device entry ports (F_Ports) to each E_Port connected to an ISL. For example, if a fabric element has two ISLs and six attached devices, the load from three devices is transferred through each ISL.

— If a single device has multiple F_Port connections to a director or switch, the switch assigns the data transfer load across multiple ISLs to maximize device availability.

- **Frame delivery order -** When directors or fabric switches calculate a new least-cost data transfer path through a fabric, routing tables immediately implement that path. This may result in Fibre Channel frames being delivered to a destination device out of order, because frames transmitted over the new (shorter) path may arrive ahead of previously-transmitted frames that traverse the old (longer) path. This causes problems because many Fibre Channel devices cannot receive frames in the incorrect order.

ATTENTION ! Activating a preferred path can result in receipt of out-of-order frames if the preferred path differs from the current path, if input and output (I/O) is active from the source port, and if congestion is present on the current path.

A rerouting delay parameter can be enabled at the Element Manager application to ensure a director or switch provides correct frame order delivery. The delay period is equal to the error detect time out value (E_D_TOV) specified in the Element Manager application. Class 2 frames transmitted into the fabric during this delay period are rejected; Class 3 frames are discarded without notification. By default, the rerouting delay parameter is disabled.

NOTE: To prevent E_Port segmentation, the same E_D_TOV and resource allocation time out value (R_A_TOV) must be specified for each fabric element.

- **E_Port segmentation** - When an ISL activates, the two fabric elements exchange operating parameters to determine if they are compatible and can join to form a single fabric. If the elements are incompatible, the connecting E_Port at each director or switch segments to prevent the creation of a single fabric. A segmented link transmits only Class F traffic; the link does not transmit Class 2 or Class 3 traffic. The following conditions cause E_Ports to segment:
  - **Incompatible operating parameters** - Either the R_A_TOV or E_D_TOV is inconsistent between the two fabric elements.
  - **Duplicate Domain_IDs** - One or more Domain_ID conflicts are detected.
  - **Incompatible zoning configurations** - Zoning configurations for the two fabric elements are not compatible. For an explanation, refer to *Configuring Zones*.
  - **Build fabric protocol error** - A protocol error is detected during the process of forming the fabric.
  - **No principal switch** - No director or switch in the fabric is capable of becoming the principal switch.

— **No response from attached switch** - After a fabric is created, each element in the fabric periodically verifies operation of all attached switches and directors. An ISL segments if a switch or director does not respond to a verification request.

— **ELP retransmission failure timeout** - A director or switch that exhibits a hardware failure or connectivity problem cannot transmit or receive Class F frames. The director or switch did not receive a response to multiple exchange link parameters (ELP) frames, did not receive a fabric login (FLOGI) frame, and cannot join an operational fabric.

- **Fabric services and state change notifications** - In a multiswitch fabric, director-provided services such as name service, registered state change notifications (RSCNs), and zoning are provided on a fabric-wide basis. For example, if a fabric-attached device queries a director or switch name server to locate all devices that support a specified protocol, the reply includes all fabric devices that support the protocol that are in the same zone as the requesting device, not just devices attached to the director or switch.

  RSCNs are transmitted to all registered device N_Ports attached to the fabric if either of the following occur:

  — A fabric-wide event occurs, such as a director logging in to the fabric, a director logging out of the fabric, or a reconfiguration because of a director or ISL failure.

  — A zoning configuration change.

- **Zoning configurations for joined fabrics** - In a multiswitch fabric, zoning is configured on a fabric-wide basis, and any change to the active zone set is applied to all directors and switches. To ensure zoning is consistent across a fabric, the following rules are enforced when two fabrics (zoned or unzoned) join through an ISL.

  — **Fabric A unzoned and Fabric B unzoned** - The fabrics join successfully, and the resulting fabric remains unzoned.

  — **Fabric A zoned and Fabric B unzoned** - The fabrics join successfully, and fabric B automatically inherits the zoning configuration from fabric A.

  — **Fabric A unzoned and Fabric B zoned** - The fabrics join successfully, and fabric A automatically inherits the zoning configuration from fabric B.

— **Fabric A zoned and Fabric B zoned** - The fabrics join successfully only if the zone sets can be merged. If the fabrics cannot join, the connecting E_Ports segment and the fabrics remain independent.

Zone sets for two directors or switches are compatible (the fabrics can join) only if the zone names for each fabric element are unique. The zone names for two fabric elements can be the same only if the zone member WWNs are identical for each duplicated zone name.

General Fabric Design Considerations

To be effective, the fabric topology design must:

- Solve the customer's business problem and provide the required level of performance.
- Meet the customer's requirements for high availability.
- Be scalable to meet future requirements.

Fabric Initialization

When multiple directors or switches are connected, E_Port (ISL) communication must be established between fabric elements and the fabric must be initialized. During fabric initialization, the fabric elements:

- Establish the operating mode for connected E_Port pairs and exchange link parameters (E_Port names, timeout values, class-specific information, and flow control parameters).
- Exchange fabric parameters, select a principal switch, and assign Domain_IDs to all switches.
- Employ a routing protocol to establish the shortest path through the fabric and program route tables for each fabric element.
- Exchange the active zone set to ensure uniform zoning is enforced between all fabric elements.

Fabric initialization is not a serial process. The process executes concurrently across all ISLs in the fabric, causing a massive flood of Class F traffic that must processed to the embedded port of each fabric element within a specified (fabric-wide) E_D_TOV.

If the fabric consists of a large number of elements (and therefore ISLs), Class F traffic may not be processed within the E_D_TOV, resulting in error recovery operations, timeouts, segmented links, or fabric failure.

Because of these problems, a fabric with a high ISL count is more difficult to build. Problems associated with a large fabric are not directly related to the large number of fabric elements but to the large number of ISLs associated with the elements.

Installing high-port count directors (such as the Intrepid 10000 Director) as SAN building blocks provides a larger number of non-blocking Fibre Channel ports. Large fabrics built around these directors require fewer additional fabric elements (smaller directors and fabric switches) and ISLs. The Intrepid 10000 Director also supports high-bandwidth (10.2000 Gbps) ISLs that reduce the fabric's total ISL count. In addition, fabrics and sub-fabrics can be merged or maintained as separate entities through flexible partitioning provided by the Intrepid 10000 Director.

Large fabrics benefit from deterministic non-blocking performance, less ISL congestion, and better cable management - performance that is not possible from a fabric constructed with smaller port-count switches interconnected with multiple ISLs.

Fabric Performance

During the design phase of a Fibre Channel fabric, performance requirements of the fabric and of component directors, fabric switches, and devices must be identified and incorporated. An effective fabric design can accommodate changes to performance requirements and incorporate additional directors, switches, devices, ISLs, and higher speed links with minimal impact to fabric operation. Performance factors that affect fabric design include:

- Application input/output (I/O) requirements, both in Gbps and I/Os per second (IOPS).
- Storage port fan-out.
- Hardware limits, including the maximum directors and switches per fabric, maximum number of ISLs per director or switch, and maximum hops between devices. For additional information, refer to *Fabric Topology Limits*.

- Software limits, including the maximum number of fabric elements managed by the SAN management application and the maximum number of zones and zone members. For additional information, refer to *SAN Management Applications* and *Configuring Zones*.

I/O Requirements

McDATA directors and fabric switches are designed with non-blocking architecture; therefore any two switch ports can communicate at the full Fibre Channel bandwidth of 1.0625, 2.1250, 4.2500, or 10.2000 Gbps without impact to other switch ports. Because most SAN-attached devices are not capable of generating I/O traffic at the full bandwidth, there is little potential for congestion between two devices attached through a single director or switch.

However, when multiple directors or switches are connected through a fabric ISL that multiplexes traffic from several devices, significant potential for congestion arises. To minimize congestion, factors such as application I/O profiles, ISL oversubscription, and device locality must be included in the fabric design.

Application I/O Profiles

Understanding application I/O characteristics is essential to SAN, fabric, and ISL design. Factors that may affect application I/O include:

- **Read/write mixture -** Although application I/O is typically a mixture of read and write operations, some applications are very biased. For example, video server applications are almost 100% read intensive, while real-time video editing applications are mostly write intensive. Read operations typically take less time than write operations, therefore storage devices for a read-intensive application usually wait for data transfer. As a consequence, read-intensive applications typically require high bandwidth to the device.
- **Type of data access -** When an application requires data, access to that data is random or sequential. For example, e-mail server activity is random access, while seismic data processing for the oil and gas industry is sequential access. Sequential data access typically takes less time than random data access, therefore sequential-access applications usually wait for data transfer. As a consequence, sequential-access applications typically require high bandwidth to the device.

- **I/O block size -** The third characteristic of application I/O is data block size, which typically ranges from two kilobytes (KB) to over one megabyte (MB). Applications that generate large blocks of data require high bandwidth to the device.

Prior to fabric design, application I/O profiles should be estimated or established that classify the application bandwidth requirements. Bandwidth consumption is classified as light, medium, or heavy. These classifications must be considered when planning ISL and device connectivity. For information about application I/O (in Gbps) and fabric performance problems due to ISL connectivity, refer to *ISL Oversubscription*. For information about application I/O (in IOPS) and fabric performance problems due to port contention, refer to *Device Fan-Out Ratio*.

ISL Oversubscription

ISL oversubscription (or congestion) occurs when multiplexed traffic from several devices is transmitted across a single ISL. When an ISL is oversubscribed, fabric elements use fairness algorithms to interleave data frames from multiple devices, thus giving fractional bandwidth to the affected devices. Although all devices are serviced, ISL and fabric performance is reduced. Figure 3-12 illustrates ISL oversubscription.

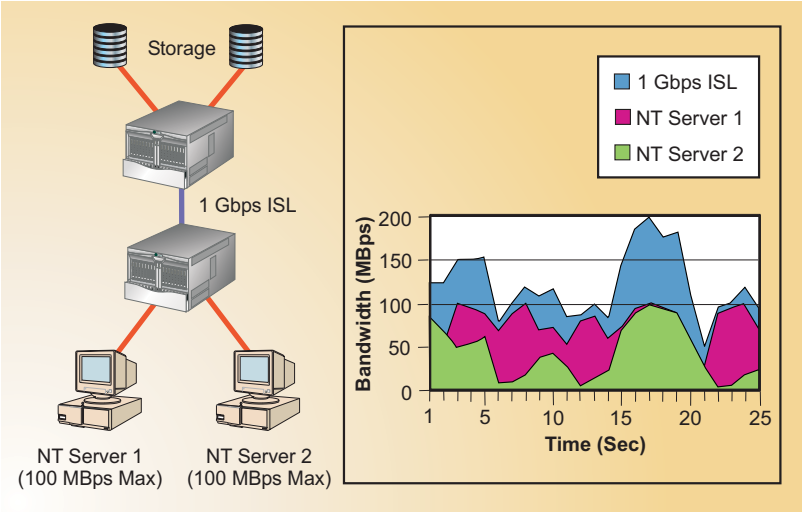

Figure 3-12 ISL Oversubscription

Two NT servers, each with maximum I/O of 100 MBps, are contending for the bandwidth of a single ISL operating at 1.0625 Gbps. In addition to data, the ISL must also transmit Class F traffic internal to the fabric. When operating at peak load, each NT server receives less than half the available ISL bandwidth.

Depending on fabric performance requirements and cost, there are several options (best practices) to solve ISL oversubscription problems, including:

- **Employ device locality -** NT Server 1 and its associated storage device can be connected through one director. NT Server 2 and its associated storage device can be connected through the other director. As a result, minimal traffic flows across the ISL between directors and the congestion problem is mitigated. For additional information, refer to *Device Locality*.
- **Install an additional ISL -** A second ISL can be installed to balance the traffic load between fabric elements. Two ISLs are sufficient to support the bandwidth of both NT servers operating at peak load.
- **Upgrade the existing ISL -** Software, firmware, and hardware can be upgraded to support a 2.1250, 4.2500, or 10.2000 Gbps bandwidth traffic load between fabric elements. A 2.1250, 4.2500, or 10.2000 Gbps ISL is sufficient to support the bandwidth of both NT servers operating at peak load.
- **Deliberately employ ISL oversubscription -** SANs are expected to function well, even with oversubscribed ISLs. Device I/O is typically bursty, few devices operate at peak load for a significant length of time, and device loads seldom peak simultaneously. As a result, ISL bandwidth is usually not fully allocated, even for an oversubscribed link. An enterprise can realize significant cost savings by deliberately designing a SAN with oversubscribed ISLs that provide connectivity for noncritical applications.

Device Locality

Devices that communicate with each other through the same director or switch have high locality. Devices that must communicate with each other through one or more ISLs have low locality. Part (A) of Figure 3-13 illustrates high device locality with little ISL traffic. Part (B) of Figure 3-13 illustrates low device locality.

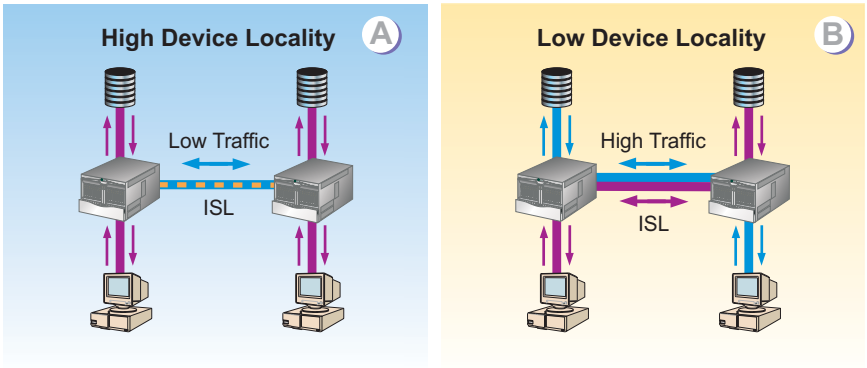

Figure 3-13 Device Locality

Although it is possible to design a SAN that delivers sufficient ISL bandwidth in a zero-locality environment, it is preferable to design local, one-to-one connectivity for heavy-bandwidth applications such as video server, seismic data processing, or medical 3D imaging.

When designing a core-to-edge fabric, servers and storage devices that support such bandwidth-intensive applications should be attached to core directors as Tier 1 devices. As a best practices policy (assuming 1.0625 Gbps ISLs), devices that generate a sustained output of 35 MBps or higher are candidates for Tier 1 connectivity. FICON devices also must use Tier 1 connectivity. For additional information, refer to *FCP and FICON in a Single Fabric*.

Device Fan-Out Ratio

The output of most host devices is bursty in nature, most devices do not sustain full-bandwidth output, and it is uncommon for the output of multiple devices to peak simultaneously. These variations are why multiple hosts can be serviced by a single storage port. This device sharing leads to the concept of fan-out ratio.

Device fan-out ratio is defined as the storage or array port IOPS divided by the attached host IOPS, rounded down to the nearest whole number. A more simplistic definition for device fan out is the ratio of host ports to a single storage port. Fan-out ratios are typically device dependent. In general, the maximum device fan-out ratio supported is 12 to 1. Figure 3-14 illustrates a fan-out ratio of 10 to 1.

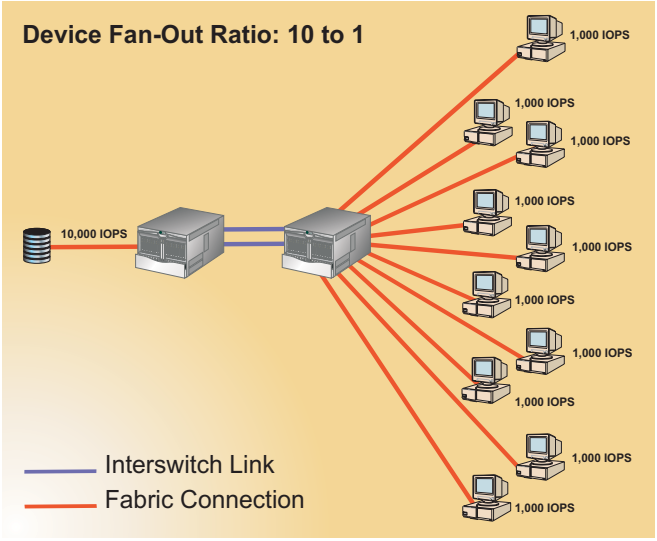

Figure 3-14 Device Fan-Out Ratio

Performance Tuning

When designing or tuning a fabric for performance, it is critical to understand application I/O characteristics so that:

- Device output in Gbps does not oversubscribe ISLs, leading to fabric congestion.
- Device output in IOPS does not result in a connectivity scheme that exceeds fan-out ratios, leading to port congestion.

Figure 3-15 illustrates performance tuning for a simple fabric using appropriate ISL connectivity, device locality, and fan-out regions for device connectivity. The fabric is comprised of one core director and six edge switches. Tier 2 servers connect to three switches at the bottom of the figure, and Tier 2 storage devices connect to three switches at the top of the figure.

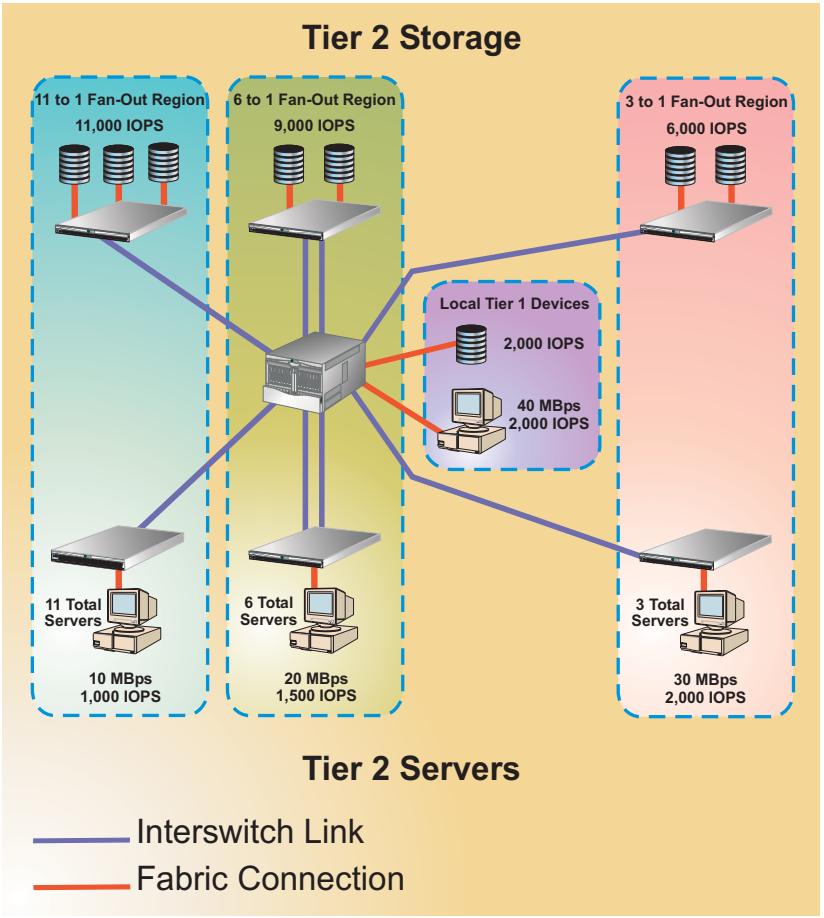

Figure 3-15 Fabric Performance Tuning

The fabric is divided into four performance regions as follows:

- **Local Tier 1 devices -** A video server application with I/O capabilities of 40 MBps and 2,000 IOPS must be connected to the fabric. Because the application is critical and high bandwidth (in excess of 35 MBps), the server and associated storage are directly attached to the core director as Tier 1 devices. No ISLs are used for server-to-storage connectivity.

- **11 to 1 fan-out region -** Eleven NT servers with I/O capabilities of 10 MBps and 1,000 IOPS are fabric-attached through a 32-port edge switch. The primary applications are e-mail and online transaction processing (OLTP). Because bandwidth use is light and noncritical, the servers are connected to the core director with a single ISL that is intentionally oversubscribed (1.1 Gbps plus Class F traffic). The servers are connected to storage devices with I/O capabilities of 11,000 IOPS.
- **6 to 1 fan-out region -** Six servers with I/O capabilities of 20 MBps and 1,500 IOPS are fabric-attached through a 16-port edge switch. Bandwidth use is light to medium but critical, so the servers are connected to the core director with two ISLs (0.6 Gbps each plus Class F traffic). The servers are connected to storage devices with I/O capabilities of 9,000 IOPS.
- **3 to 1 fan-out region -** Three servers with I/O capabilities of 30 MBps and 2,000 IOPS are fabric-attached through a 16-port edge switch. Bandwidth use is medium but non critical, so the servers are connected to the core director with one ISL (0.9 Gbps plus Class F traffic). The servers are connected to storage devices with I/O capabilities of 6,000 IOPS.

Fabric Availability

Many fabric-attached devices require highly-available connectivity to support applications such as disk mirroring, server clustering, or business continuance operations. High availability is accomplished by deploying a resilient fabric topology or redundant fabrics.

A fabric topology that provides at least two internal routes between fabric elements is considered resilient. A single director, switch, or ISL failure does not affect the remaining elements and the overall fabric remains operational. However, unforeseen events such as human error, software failure, or disaster can cause the failure of a single resilient fabric. Using redundant fabrics (with resiliency) mitigates these effects and significantly increases fabric availability.

Fibre Channel fabrics are classified by four levels of resiliency and redundancy. From least available to most available, the classification levels are:

- **Nonresilient single fabric -** Directors and switches are connected to form a single fabric that contains at least one single point of failure (fabric element or ISL). Such a failure causes the fabric to fail and segment into two or more smaller fabrics.

- **Resilient single fabric -** Directors and switches are connected to form a single fabric, but no single point of failure can cause the fabric to fail and segment into two or more smaller fabrics.
- **Nonresilient dual fabric -** Half of the directors and switches are connected to form one fabric, and the remaining half are connected to form an identical but separate fabric. Servers and storage devices are connected to both fabrics. Each fabric contains at least one single point of failure (fabric element or ISL). All applications remain available, even if an entire fabric fails.
- **Resilient dual fabric -** Half of the directors and switches are connected to form one fabric, and the remaining half are connected to form an identical but separate fabric. Servers and storage devices are connected to both fabrics. No single point of failure can cause either fabric to fail and segment. All applications remain available, even if an entire fabric fails and elements in the second fabric fail.

  A dual-fabric resilient topology is generally the best design to meet high-availability requirements. Another benefit of the design is the ability to proactively take one fabric offline for maintenance or upgrade without disrupting SAN operations.

Redundant Fabrics

If high availability is important enough to require dual-connected servers and storage, a dual-fabric solution is generally preferable to a dual-connected single fabric. Dual fabrics maintain simplicity and reduce (by 50%) the size of fabric routing tables, name server tables, updates, and Class F management traffic. In addition, smaller fabrics are easier to analyze for performance, fault isolate, and maintain.

Figure 3-16 illustrates simple redundant fabrics. Fabric “A” and fabric “B” are symmetrical, each containing one core director and four edge switches. All servers and storage devices are connected to both fabrics.

Some dual-attached devices support active-active paths, while others support only active-passive paths. Active-active devices use either output path equally, and thus use both fabrics and double the device bandwidth. Active-passive devices use the passive path only when the active path fails.

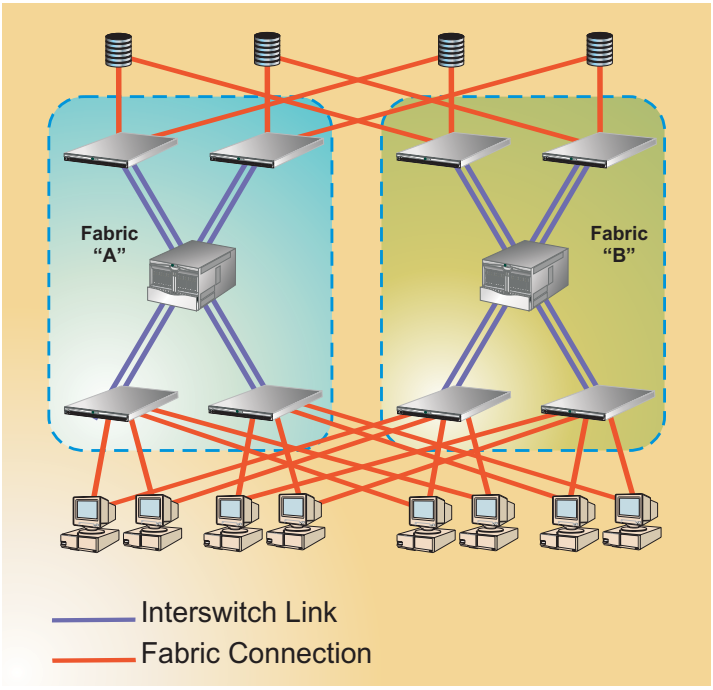

Figure 3-16 Redundant Fabrics

When deploying redundant fabrics, it is not required that the fabrics be symmetrical. As an example, single-attached devices, such as tape drives and noncritical servers and storage, can be logically grouped and attached to one of the fabrics.

Fabric Scalability

Businesses are experiencing an unprecedented growth of information and the requirement to maintain that information online. To meet these requirements, Fibre Channel SANs provide the theoretical infrastructure to connect thousands of servers to hundreds of storage devices. To provide enterprise-class performance, scalable fabric designs are required. Refer to *Chapter 4, Implementing SAN Internetworking Solutions* for detailed information.

A scalable fabric allows for nondisruptive addition of fabric elements (directors, fabric switches, and SAN routers) or ISLs to increase the size or performance of the fabric or SAN. Large, scalable fabrics and SANs are constructed by incorporating:

- **High-port count directors -** Installing high-port count directors as SAN building blocks provides a larger number of non-blocking Fibre Channel ports per fabric element and reduces the need for ISLs. Newer products support high-bandwidth (10.2000 Gbps) ISLs that also reduce the fabric ISL count. In addition, fabrics and sub-fabrics can be merged or maintained as separate entities through dynamic partitioning. Refer to *General Fabric Design Considerations* and *Inter-FlexPar Routing* for information.
- **SAN routers -** Installing SAN routers provides interoperable E_Port connectivity between local SAN fabrics. However, SAN routers terminate the E_Port connection at each SAN edge. This allows devices in each SAN to communicate through the router, but preserves the autonomy of each local SAN. Refer to *R_Port Operation* and *Inter-FlexPar Routing* for information.

Scalability also relates to investment protection. If a core fabric switch is replaced with a newer or higher port count switch (such as the Intrepid 10000 Director), it is often valuable to use the existing switch elsewhere in the fabric (at the edge).

Obtaining Professional Services

Planning and implementing a multiswitch fabric topology can be a complex and difficult task. Obtain planning assistance from McDATA's professional services organization before implementing a fabric topology.

Mixed Fabric Design Considerations

This section discusses mixed fabric design considerations, including:

- Fibre Channel Protocol (FCP) and FICON environments in a single fabric.
- Multiple data transmission speeds (1.0625, 2.1250, 4.2500, and 10.2000 Gbps) in a single fabric.

FCP and FICON in a Single Fabric

Fibre Channel Layer 4 (FC-4) describes the interface between Fibre Channel and various upper-level protocols. FCP and FICON are the major FC-4 protocols. FCP is the Fibre Channel protocol that supports the small computer system interface (SCSI) upper-level transport protocol. FICON is the successor to the enterprise systems connection (ESCON) protocol and adds increased reliability and integrity to that provided by the FCP protocol.

Because FCP and FICON are both FC-4 protocols, routing of Fibre Channel frames is not affected when the protocols are mixed in a single fabric environment. However, management differences in the protocols arise when a user changes director or fabric switch parameters through zoning or connectivity control. In particular:

- FCP communication parameters are port number and name-centric, discovery oriented, assigned by the fabric, and use the Fibre Channel name server to control device communication.
- FICON communication parameters are logical port address-centric, definition oriented, assigned by the attached host, and use host assignment to control device communication.

Considerations that need to be evaluated when intermixing FCP and FICON protocols are:

- Director or switch management.
- Port numbering versus port addressing.
- Management limitations.
- Features that impact protocol intermixing.
- Best practices.

Director or Switch Management

When intermixing FCP and FICON protocols, it must be determined if the director or fabric switch is to be operated using the open systems or FICON management style. This setting only affects the operating mode used to manage the director or switch; it does not affect F_Port operation. FCP devices can communicate with each other when the attached fabric element is set to FICON management style, and FICON devices can communicate with each other when the attached fabric element is set to open systems management style.

- When a director or fabric switch is set to open systems management style, FCP connectivity is defined within a Fibre Channel fabric using WWNs of devices that are allowed to form connections. When connecting to the fabric, an FCP device queries the name server for a list of devices to which connectivity is allowed. This connectivity is hardware-enforced through a name server zoning feature that partitions attached devices into restricted-access zones.
- When a director or fabric switch is set to FICON management style, host-to-storage FICON connectivity and channel paths are defined by a host-based hardware configuration definition (HCD) program and a director or switch-resident management server called the control unit port (CUP). A user-configured (director or switch-resident) prohibit dynamic connectivity mask (PDCM) array associates logical port addresses. FICON devices do not query the name server for accessible devices because connectivity is defined at the host, and additionally at the director or switch. This connectivity is hardware-enforced in the routing tables of each port.

NOTE: The Intrepid 10000 Director and Sphereon 4300, 4400, and 4500 Fabric Switches do not support operation using FICON management style nor transmission of FICON frames.

PDCM connectivity control is configured and managed at the director or switch level using the *Configure Allow/Prohibit Matrix - Active* dialog box (Figure 5-4). For additional information, refer to *PDCM Arrays*.

ATTENTION ! When configuring a PDCM array that prohibits E_Port connectivity, mistakes can render ISLs unusable and cause complex routing problems. These problems can be difficult to fault isolate and sometimes manifest incorrectly as end-device issues.

Port Numbering Versus Port Addressing

Consideration must be given to the implications of port numbering for the FCP protocol versus logical port addressing for the FICON protocol. FCP configuration attributes are implemented through zoning. Zones are configured through the associated Element Manager application by:

- The eight-byte (64-digit) WWN assigned to the host bus adapter (HBA) or Fibre Channel interface installed in a device connected to the director or switch.

- The Domain_ID and physical port number of the director or fabric switch port to which a device is attached.

FICON configuration attributes are implemented through logical port addressing. This concept is consistent with the address-centric nature of FICON and allows ports to be swapped for maintenance operations without regenerating a host configuration. For McDATA products, logical port addresses are derived by converting the port number from numerical to hexadecimal format and adding a hexadecimal four to the result. Figure 3-17 illustrates port numbering and logical port addressing for Intrepid 6140 Director ports accessed from the front.

UPM Cards								CTP - 1 Card	CTP - 0 Card	UPM Cards							
31	30	29	28	27	26	25	24			23	22	21	20	19	18	17	16
127 7F **83**	123 7B **7F**	119 77 **7B**	115 73 **77**	111 6F **73**	107 6B **6F**	103 67 **6B**	99 63 **67**			95 5F **63**	91 5B **5F**	87 57 **5B**	83 53 **57**	79 4F **53**	75 4B **4F**	71 47 **4B**	67 43 **47**
126 7E **82**	122 7A **7E**	118 76 **7A**	114 72 **76**	110 6E **72**	106 6A **6E**	102 66 **6A**	98 62 **66**			94 5E **62**	90 5A **5E**	86 56 **5A**	82 52 **56**	78 4E **52**	74 4A **4E**	70 46 **4A**	66 42 **46**
125 7D **81**	121 79 **7D**	117 75 **79**	113 71 **75**	109 6D **71**	105 69 **6D**	101 65 **69**	97 61 **65**			93 5D **61**	89 59 **5D**	85 55 **59**	81 51 **55**	77 4D **51**	73 49 **4D**	69 45 **49**	65 41 **45**
124 7C **80**	120 78 **7C**	116 74 **78**	112 70 **74**	108 6C **70**	104 68 **6C**	100 64 **68**	96 60 **64**			92 5C **60**	88 58 **5C**	84 54 **58**	80 50 **54**	76 4C **50**	72 48 **4C**	68 44 **48**	64 40 **44**
60 3C **40**	56 38 **3C**	52 34 **38**	48 30 **34**	44 2C **30**	40 28 **2C**	36 24 **28**	32 20 **24**			28 1C **20**	24 18 **1C**	20 14 **18**	16 10 **14**	12 0C **10**	08 08 **0C**	04 04 **08**	00 00 **04**
61 3D **41**	57 39 **3D**	53 35 **39**	49 31 **35**	45 2D **31**	41 29 **2D**	37 25 **29**	33 21 **25**			29 1D **21**	25 19 **1D**	21 15 **19**	17 11 **15**	13 0D **11**	09 09 **0D**	05 05 **09**	01 01 **05**
62 3E **42**	58 3A **3E**	54 36 **3A**	50 32 **36**	46 2E **32**	42 2A **2E**	38 26 **2A**	34 22 **26**			30 1E **22**	26 1A **1E**	22 16 **1A**	18 12 **16**	14 0E **12**	10 0A **0E**	06 06 **0A**	02 02 **06**
63 3F **43**	59 3B **3F**	55 37 **3B**	51 33 **37**	47 2F **33**	43 2B **2F**	39 27 **2B**	35 23 **27**			31 1F **23**	27 1B **1F**	23 17 **1B**	19 13 **17**	15 0F **13**	11 0B **0F**	07 07 **0B**	03 03 **07**
15	14	13	12	11	10	9	8			7	6	5	4	3	2	1	0

Figure 3-17 Intrepid 6140 Port Numbers and Logical Port Addresses (Front)

The figure shows:

- CTP card positions (**0** and **1**).
- Universal port module (UPM) card numbers at the top and bottom (numerical **0** through **31**).
- Numerical physical port numbers in blue (**00** through **127**).
- Hexadecimal physical port numbers in red (**00** through **7F**).
- Logical port addresses in bold (hexadecimal **04** through **83**).

Figure 3-18 illustrates port numbering and logical port addressing for Intrepid 6140 Director ports accessed from the rear.

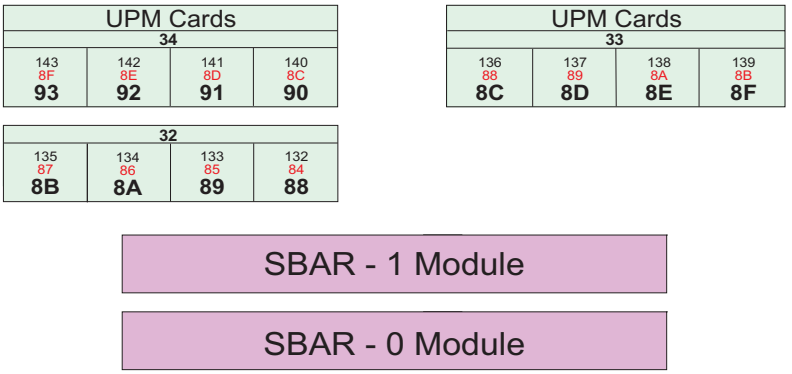

Figure 3-18 Intrepid 6140 Port Numbers and Logical Port Addresses (Rear)

The figure shows:

- SBAR positions (**0** and **1**).
- UPM card numbers (numerical **32**, **33**, and **34**).
- Numerical physical port numbers in blue (**132** through **143**). Port numbers **128** through **131** are embedded and not addressable.
- Hexadecimal physical port numbers in red (**84** through **8F**). Port numbers **80** through **83** are embedded and not addressable.
- Logical port addresses in bold (hexadecimal **88** through **93**). Port addresses **84** through **87** are not addressable.

Although Figure 3-17 and Figure 3-18 depict UPM card maps only for the Intrepid 6140 Director, physical port numbers and logical port addresses can be extrapolated for the Intrepid 6064 Director (64 ports), Sphereon 3232 Fabric Switch (32 ports), Sphereon 4300 Fabric Switch (12 ports), Sphereon 4400 Fabric Switch (16 ports), Sphereon 4500 Fabric Switch (24 ports), and Sphereon 4700 Fabric Switch (32 ports).

Management Limitations

The following considerations must be given to the limitations and interactions of director or fabric switch management when using open systems (FCP) or FICON management style:

- FICON port-to-port connectivity is hardware enforced, while FCP port-to-port connectivity is software or hardware enforced (depending on the director or switch firmware release level).
  - — FICON architecture controls connectivity through a host-based HCD program, the CUP, and a director or switch-resident PDCM array. The CUP and PDCM array support hardware enforcement of connectivity control to all port connections; therefore when a director or switch is set to FICON management style, zoning information is restricted by the hardware instead of by the name server.
  - — When a director or switch is set to open systems management style, CUP support and the PDCM array are disabled. For FICON devices attached to the director or switch, the user must manage connectivity to match logical port addressing established through the host-based HCD program. For example if a FICON hosts expects connectivity through logical port address **1C**, the user must ensure the host is connected to physical port number **24**. Refer to Figure 3-17 and Figure 3-18 for physical port number and logical port address maps.
- The FCP protocol supports multiple domains (multiswitch fabrics). The FICON protocol may or may not be limited to a single domain (single-switch fabric), depending on the director or switch firmware release level as follows:
  - — For earlier versions of director or switch firmware (prior to Enterprise Operating System, classic (E/OSc) Version 4.0), the FICON protocol is limited to a single domain (single-switch fabric) due to single-byte Fibre Channel link address limitations inherited from ESCON. Consequently, when a director or switch is set to FICON management style (FICON compliant), E_Port connections (ISLs) are not allowed with another fabric switch. The director or switch reports an attempted E_Port connection as invalid and prevents the port from coming online.
  - — For later versions of director or switch firmware (E/OSc Version 4.0 and later), the domain field of the destination ID is added to the Fibre Channel link address, thus specifying the link address on source and target fabric elements and enabling E_Port (ISL) connectivity. This connectivity is called FICON cascading. For additional information, refer to *FICON Cascading*.

- When employing inband (Fibre Channel) director or switch management, the open-systems management server (OSMS) is associated with the FCP protocol, and the FICON management server (FMS) is associated with the FICON protocol. Management server differences tend to complicate security and control issues.

  NOTE: The Intrepid 10000 Director and Sphereon 4300, 4400, and 4500 Fabric Switches do not support out-of-band management through FMS.

  Each server provides facilities to change zoning information (FCP protocol) or the logical port address-based connectivity configuration (FICON protocol), but neither provides sufficient functionality for both protocols.

Features that Impact Protocol Intermixing

McDATA supports the following features that impact how a director or switch behaves when deployed in an intermixed environment:

- Hardware-enforced zoning.
- SANtegrity Binding (including fabric and switch binding).
- FICON cascading.

Hardware-Enforced Zoning

Hardware-enforced zoning (hard zoning) allows a user to program director or switch route tables that enable hardware logic to route Fibre Channel frames. This process prevents traffic between source and destination devices not in the same zone. Hard zoning provides the open-systems environment with the same protection that PDCM arrays provide in the FICON environment.

In environments that include discovery-oriented devices (FCP) and definition-oriented devices (FICON), system administrators must keep device definitions and zoning definitions synchronized. Hard zoning enforces zoning information at the director or switch level and ensures the information takes precedence over access definitions configured at the device level. This provides a security element that is useful for mixed environments using both definition and discovery. For additional information, refer to *Zoning*.

SANtegrity Binding

McDATA offers a SANtegrity Binding feature (including both fabric binding and switch binding) that allows the creation of reliable SAN configurations and provides a mechanism for attached devices to query the user-configured security level employed in a SAN. The feature significantly reduces the impacts of accidental or operator-induced errors.

Fabric binding defines the directors and switches allowed to participate in a fabric, thus preventing accidental fabric merges. Switch binding defines the devices allowed to connect to directors and switches in a fabric, thus providing additional security in SAN environments that must manage a large number of devices. For additional information, refer to *SANtegrity Binding*.

FICON Cascading

FICON is most often deployed in SANs that have high data integrity and reliability standards. However, the initial FICON architecture was limited to one domain (i.e. a single-switch fabric), which creates severe distance and connectivity limitations. These data standards and the requirement for FICON fabrics in SANs led to protocol changes that support FICON cascading.

FICON cascading allows an IBM eServer zSeries processor to communicate with other zSeries processors or peripheral devices (such as disks, tape libraries, or printers) through a fabric consisting of two or more FICON directors or switches. Cascaded FICON fabrics also provide high end-to-end data integrity to ensure changes to a data stream are always detected and rectified and that data is always delivered to the correct fabric end point. For additional information, refer to *FICON Cascading*.

A related feature to consider is the announcement of FCP support for IBM eServer zSeries processors. This development accelerates the requirement for intermix protocol fabrics, since primary processors now support both FICON and FCP.

Protocol Intermixing Best Practices

The Element Manager graphical user interface (GUI) provides an open systems or FICON management style. Users can toggle between management styles with the director or switch online.

However, the firmware and SAN management applications do not prevent FCP and FICON device configurations that may interfere with each other. A successful intermix environment requires a set of best practice conventions as follows:

1. **Upgrade fabric element firmware to a common version -** Ensure fabric elements are operating at a common firmware level. This reduces errors due to director or switch incompatibility. E/OSc Version 4.0 or higher is required to support FICON cascading. E/OSc Version 6.0 or higher is recommended.
2. **Upgrade fabric element software to a common version -** Ensure fabric elements are operating at a common software level. This simplifies fabric fault isolation and reduces errors due to director or switch incompatibility. SANavigator 4.0 or EFCM 8.0 (or higher) is required to support a unified management style and is recommended.

   — When a director or switch is set to open systems management style, a traditional Fibre Channel fabric is supported. Inband management through the FMS or OSMS is also supported. The key concern is to avoid disrupting installed FCP devices when connecting FICON devices to a fabric element and modifying configurations to facilitate FICON communication. The Element Manager application does not use logical port addressing or display the *Configure Allow/Prohibit Matrix - Active* dialog box. A PDCM array is not supported, and the HCD defined by an attached host describes FICON connectivity requirements.

   — When a director or switch is set to FICON management style, either multiple domains (fabric elements) are supported, or only a single domain (fabric element) is supported, depending on the firmware release level. Inband management through the FMS or OSMS is also supported. The Element Manager application provides a PDCM array configured at the *Configure Allow/ Prohibit Matrix - Active* dialog box. The array activates all or a subset of the connectivity paths established by a host-based HCD.

   When using firmware prior to E/OSc Version 4.0 and the FICON management style, ports are set to F_Port operation, thus eliminating E_Port capability (ISL and fabric capability).

— When using inband director or switch management, either (or both) of the FMS or OSMS features can be enabled. When either (or both) features are enabled, the director or switch can be set to open systems or FICON management style.

3. **Upgrade fabric elements to a common feature set -** Ensure a common set of PFE-keyed optional features (refer to *Optional Feature Keys*) is installed on each fabric element. This reduces errors due to director or switch incompatibility. In addition, the SANtegrity Binding feature (with *Enterprise Fabric Mode* enabled) is required to support FICON cascading.

4. **Logically assign ports -** To organize devices into manageable groups for zoning, director or switch ports should be logically assigned to FCP port groups and FICON port groups. Although FICON devices can be zoned by device WWN, they must also be assigned logical port addresses that correspond to the port addresses configured by the attached host HCD. FICON devices must be attached to these assigned ports. In addition, PDCM arrays affect port connections at the hardware level, so a range of port addresses must be established for FCP device use, and a separate range of port addresses must be established for FICON device use. FCP ports should always be configured to allow communication with each other but disallow communication with FICON ports, and vice versa.

5. **Configure FICON cascading -** Configure and enable FICON cascading for all fabric elements. Refer to *FICON Cascading Best Practices* for instructions. As part of this step, ensure the SANtegrity Binding feature key is installed and *Enterprise Fabric Mode* enabled for all directors and switches.

— In conjunction with the SANtegrity Binding feature (fabric and switch binding), consider enabling port binding from a director or switch's Element Manager application. Port binding explicitly defines (by WWN or nickname) the device allowed to attach to a Fibre Channel port and provides additional security when logically allocating ports to FCP and FICON groups. Although this process creates additional configuration overhead, port binding is useful for implementations that require protection from accidental misconfigurations.

6. **Configure PDCM arrays -** For each director or switch managed by the FICON management style, define the allow and prohibit settings for FICON device connectivity. Use the Element Manager application's *Configure Allow/Prohibit Matrix - Active* dialog box. Port connectivity assignment (step 4) should be reflected in PDCM arrays for FICON connectivity management. The baseline configuration for each fabric element must prohibit communication between FICON and FCP devices.

   — Because PDCM arrays affect port connections at the hardware level, it is imperative to establish a range of port addresses for FCP use and another range for FICON use. FCP-assigned ports should be configured to allow communication with each other and prohibit communication with FICON-assigned ports, and vice versa.

   — At the *Configure Allow/Prohibit Matrix - Active* dialog box, assigning port names to logical port addresses is another practice that should be followed. For example, the port name for all FCP devices could begin with FCP or OS to indicate the associated port addresses attach to open-systems devices. This information emphasizes which ports are FCP ports and which are FICON ports and gives a user the ability to better manage the connectivity matrix.

   — Caution should be exercised when using a PDCM array to prohibit E_Port connectivity. For additional information, refer to *PDCM Arrays*.

7. **Configure zoning -** Well-behaved intermix environments require the creation of separate zones for FCP and FICON devices. Group all FICON devices into one zone, then group FCP devices into multiple zones in traditional fashion to facilitate typical open-systems communication.

   — Be aware that FICON devices do not use the Fibre Channel name server, therefore name server-based zoning does not affect FICON connectivity. However, the name server does affect distribution of registered state change notification (RSCN) service requests to FICON devices. If a FICON device is not in the same zone as other devices, state changes are not properly communicated.

— All FICON devices must be included in the same zone to facilitate proper state change notification. This is achieved by creating a unique FICON zone or using the default zone. Disable the default zone and explicitly create a unique zone for all FICON devices. Regardless of the director or switch operating mode, FCP devices must be zoned in the traditional fashion, and FICON devices must be zoned to provide isolation from the FCP devices. All FICON devices must be included in the same zone to facilitate proper state change communication.

— When establishing a zoning configuration, FICON devices must be assigned to director or switch port addresses that correspond to port HCD-assigned address definitions configured by the attached host. Associated FICON devices must be connected to the ports as configured.

— Note the reciprocal nature of zoning configurations and PDCM arrays. When configuring zoning, all FICON devices are placed in one zone and FCP devices are zoned normally. When configuring definitions in a PDCM array, all FCP devices are configured to allow communication only with each other and FICON devices are configured normally.

— FICON port addressing provides the ability to swap ports for maintenance. In general, swapping ports in intermix environments does not affect the practices described. However, if a user implements zoning using a Domain_ID and port numbers, zoning information must be updated contiguous with the port swap operation.

Multiple Data Transmission Speeds in a Single Fabric

The Sphereon 3232 Fabric Switch, Intrepid 6000-series Directors, and Sphereon 4000-series Fabric Switches support auto-sensing of 1.0625, 2.1250, and 4.2500 Gbps device connections. The Intrepid 10000 Director supports 1.0625, 2.1250, and 10.2000 Gbps connections. The introduction of a higher data transmission speed to the SAN design provides several benefits and alternatives:

- **High-speed device connectivity -** As Fibre Channel devices and HBAs evolve and become 10.2000 Gbps-capable, higher-speed switches are required to provide basic fabric connectivity.

- **Better fabric performance -** As a connection between fabric switches, a 10.2000 Gbps ISL delivers significantly greater bandwidth. Fibre Channel devices that are not 10.2000 Gbps-capable benefit from a higher-speed ISL, because slower traffic is multiplexed and transmitted through the 10.2000 Gbps ISL.
- **Additional port count -** If additional ISL bandwidth is not required for fabric performance, 10.2000 Gbps connectivity allows the number of ISL connections to be reduced, thus yielding additional director or switch ports for device connectivity.

When installing 10.2000 Gbps-capable fabric elements in a core-to-edge topology, deploy the directors at the fabric core to provide end-to-end high-speed ISL capability. If 10.2000 Gbps device connectivity is required, attach the devices to the core director as Tier 1 devices. If possible, employ device locality by connecting 10.2000 Gbps devices to the same director.

FICON Cascading

The initial FICON architecture did not permit connection of multiple directors or switches because the protocol specified a single byte for the link (port) address definition in the input-output configuration program (IOCP). The link address only defined the Port_ID for a unique domain (director or switch).

The current FICON architecture provides two-byte addressing that allows the IOCP to specify link (port) addresses for any number of domains by including the domain address with the Port_ID. FICON fabrics can now be configured using multiple directors and switches (FICON cascading). In a cascaded FICON environment, at least three Fibre Channel links are involved:

- The first link is between the FICON channel card (N_Port) of an IBM eServer zSeries processor and a director or switch F_Port.
- The second link is an ISL between two director or switch E_Ports.
- The final link is from a director or switch F_Port to a FICON adapter card (control unit N_Port) in a storage device, tape device, or other peripheral.

These Fibre Channel links connect FICON fabric elements and provide a physical transmission path between a channel and control unit. Users may configure multiple ISLs between cascaded FICON directors or switches to ensure redundancy and adequate bandwidth.

High-Integrity Fabrics

Cascaded FICON directors and switches must support high-integrity fabrics. McDATA fabric elements must have the SANtegrity Binding feature installed and operational with *Enterprise Fabric Mode* enabled. High-integrity fabric architecture support includes:

- **Fabric binding -** Only directors or switches with fabric binding installed are allowed to attach to specified fabrics in a SAN. Specifically:
  - — Fabric elements *without* a SANtegrity Binding feature key are prohibited from connecting to fabric elements *with* an active SANtegrity Binding feature key.
  - — Inherent to directors and switches with an active SANtegrity Binding feature key is a fabric membership list (comprised of acceptable WWNs and Domain_IDs) of the elements logged into the fabric. This membership list is exchanged between fabric elements, and an element with an incompatible list is isolated from the fabric. Membership list data eliminates duplicate Domain_IDs and other address conflicts and ensures a consistent, unified behavior across the fabric.
- **Switch binding -** Switch binding allows only specified devices and fabric elements to connect to specified director or switch ports.
- **Insistent Domain_ID -** When enabled through the *Enterprise Fabric Mode* dialog box, this parameter ensures duplicate Domain_IDs are not used within a fabric. It also ensures a fabric element cannot automatically change its Domain_ID when a director or switch with a duplicate Domain_ID attempts to join the fabric. The invalid (duplicate Domain_ID) fabric element is rejected, and intentional user intervention is required to change the Domain_ID to a valid number.

For additional information about the SANtegrity Binding feature, refer to *SANtegrity Binding*.

Minimum Requirements

The following are minimum hardware, firmware, and software requirements to configure and enable a FICON-cascaded SAN:

- A single-vendor switching environment with two or more of the following McDATA directors or switches:

— Intrepid 6064 or 6140 Director.

— Sphereon 3232 or 4500 Fabric Switch.

- E/OSc Version 4.0 (or later) must be installed on all directors or switches. E/OS firmware Version 6.0c (or later) is recommended. All fabric elements must be at the same firmware version level.
- The SANtegrity Binding feature key must be installed and enabled on all directors and switches. *Enterprise Fabric Mode* must also be enabled on all fabric elements.
- Enterprise Fabric Connectivity manager (EFCM) Version 6.0 or later must be installed on the product management server.
- One or more of the following IBM servers with FICON or FICON Express™ channel adapter cards:

  — eServer zSeries 800 (z800) processor.

  — eServer zSeries 900 (z900) processor.

  — eServer zSeries 990 (z990) processor.

  NOTE: FICON cascading is not supported for IBM S/390 Parallel Enterprise Servers (Generation 5 or Generation 6).

- The z/OS Version 1.3 or Version 1.4 operating system (with service as defined in PSP Buckets for device type 2032, 2042, 2064, or 2066) must be installed on the IBM server.
- Licensed Internal Code (LIC) driver 3G at microcode level (MCL) J11206 or later must be installed on the IBM server.

FICON Cascading Best Practices

A successful FICON-cascaded SAN environment requires a set of best practice conventions as follows:

1. **Connect fabric elements -** Establish one or more ISLs between cascaded directors of fabric switches as follows:

   a. Ensure fabric elements are defined to the SAN management application. If the elements must be defined, refer to the appropriate switch or director installation manual for instructions.

   b. Ensure the preferred Domain_ID for each director or switch is unique and does not conflict with the ID of another fabric element.

c. Ensure the R_A_TOV and E_D_TOV values for fabric elements are identical.

d. Route multimode or singlemode fiber-optic cables (depending on the type of transceiver installed) between customer-specified E_Ports at each fabric element.

2. **Verify operation of local FICON applications -** Ensure the ISL connection(s) do not disrupt fabric element operation nor disrupt local FICON (non-cascaded) traffic. Perform this step at each director or switch.

   a. At the SAN management application's physical map, right-click the director or switch product icon, then select *Element Manager* from the pop-up menu. The Element Manager application opens.

   b. If required, click the *Hardware* tab. The *Hardware View* (Figure 2-5) displays. Verify the status bar at the bottom left corner of the window displays a green circle, indicating director or switch status is operational. If a problem is indicated, go to *MAP 0000: Start MAP* in the product-specific *Installation and Service Manual.*

   c. Have the customer verify operation of non-cascaded FICON applications at each director or switch.

3. **Verify ISL operation -** Ensure ISL connectivity between fabric elements. Perform this step at each director or switch.

   a. At the Element Manager application's *Hardware View*, double-click the graphical E_Port connector used for the ISL. The *Port Properties* dialog box displays.

   b. Ensure the *Link Incident* field displays **None** and the *Reason* field is blank. If an ISL segmentation or other problem is indicated, go to *MAP 0000: Start MAP* in the product-specific *Installation and Service Manual.*

   c. Click *Close* to close the dialog box and return to the *Hardware View.*

4. **Install SANtegrity Binding on fabric elements -** Configure and enable the SANtegrity Binding feature at each director or switch as follows:

   a. At the Element Manager application, install the SANtegrity Binding PFE key. Refer to installation instructions in the product-specific *Installation and Service Manual*.

b. At the SAN management application, configure fabric binding. Refer to installation instructions in the *SANavigator Software Release 4.2 User Manual* (621-000013) or the *EFC Manager Software Release 8.7 User Manual* (620-000170).

c. At the Element Manager application, configure switch binding. Refer to installation instructions in the product-specific *Installation and Service Manual*.

5. **Ensure FICON devices are logged in -** Verify FICON devices are logged in to each director or switch as follows:

   a. At the Element Manager application's *Hardware View*, click the *Node List* tab. The *Node List View* displays.

   b. Inspect the node descriptors and verify the correct FICON devices (channels and control units) are logged in to each director or switch.

6. **Enable Enterprise Fabric Mode -** Enable *Enterprise Fabric Mode* as follows:

   a. Minimize the Element Manager application to display the SAN management application, then select *Enterprise Fabric Mode* from the *Configure* menu. The *Enterprise Fabric Mode* dialog box displays.

   b. Select the fabric to be configured from the *Fabric Name* drop-down list. The selected fabric's status displays in the *Enterprise Fabric Mode* field.

   c. Click *Activate* to close the dialog box and enable *Enterprise Fabric Mode* for the selected fabric.

7. **Verify FICON devices are still logged in -** Maximize the Element Manager application. Inspect the *Node List View* and verify FICON devices (channels and control units inspected in step 5) are still logged in to each director or switch.

8. **Change switch binding enforcement if required -** If the SAN environment is volatile (characterized by a high volume of optical cable connects, disconnects, and movement), change switch binding enforcement to restrict E_Ports only.

   a. At the Element Manager application, click the *Hardware* tab. At the *Hardware View*, select *Switch Binding*, then *Change State* from the *Configure* menu. The *Switch Binding - State Change* dialog box displays.

b. Ensure the *Enable Switch Binding* check box is checked (enabled).

c. Select (click) the **Restrict E_Ports** radio button to restrict connections from specific fabric elements to E_Ports. WWNs can be added to the membership list to allow fabric element connection and removed from the list to prohibit fabric element connection. Devices are allowed to connect to any F_Port or FL_Port without restriction.

d. Click *Activate* to close the dialog box and enforce the connection policy.

9. **Update channel path and control unit definitions -** A cascaded FICON environment requires channel entry switch and link address updates to the input/output configuration program (IOCP) as follows:

   a. In the IOCP, define an entry switch ID in the **SWITCH** keyword of the channel path identifier (CHPID) definition.

   NOTE: An entry switch is a fabric director or switch connected to the FICON channel of a zSeries processor and a second fabric director or switch.

   b. In the IOCP, define a 2-byte link address (consisting of a switch (or domain) address and port address) for the cascaded switch in the **LINK** keyword of the control unit (**CNTLUNIT**) definition.

   NOTE: An cascaded switch is a fabric director or switch connected to a destination control unit and an entry switch.

   c. Run the IOCP to create an input/output configuration data set (IOCDS). The switch ID (CHPID macroinstructions) and 2-byte link address (control unit macroinstructions) are updated in the IOCDS.

   Refer to the IBM *FICON Native Implementation and Reference Guide* (SG24-6266) for additional information.

10. **Verify FICON devices log back in -** Inspect the *Node List View* and verify FICON devices (channels and control units inspected in step 5) log back in to the fabric as expected.

11. **Verify cascaded FICON operation -** Have the customer verify operation of established logical FICON paths between channels and control units, and verify that cascaded FICON traffic is transmitted through the fabric as expected.

4

Implementing SAN Internetworking Solutions

Enterprise-level information technology (IT) departments often deploy storage configurations that include direct-attached storage, network-attached storage (NAS), and small, isolated storage area networks (SANs). These configurations often result in:

- Isolated and inefficiently-used applications, data storage, and computing resources.
- Costly, decentralized, and complex asset management.
- Slow transactions and data access.
- Inability to comply with government regulations that dictate data retention and security policies.

The solution for these problems is to implement a internetworking strategy that consolidates IT resources and deploys an enterprise-wide fabric. This chapter describes planning considerations for implementing SAN internetworking solutions using McDATA switch products. The chapter specifically describes:

- SAN island consolidation.
- Implementing business continuance and disaster recovery (BC / DR) solutions.
- Consolidating and integrating Internet small computer systems interface (iSCSI) servers and storage.

SAN Island Consolidation

SAN islands tend to be constructed along application (such as product test, finance, or engineering), operating system (OS), protocol, or geographical (site-based) boundaries. Because of application and OS segmentation, large data centers at single sites often consist of SAN islands constructed with relatively small Fibre Channel switches.

SAN Island Benefits

A SAN island deployment strategy provides many benefits and may be sufficient to meet an enterprise's needs because:

- SAN islands serve a purpose. The enterprise buys a switch, builds a simple fabric, and implements a SAN around a particular application.
- Data, applications, and operating systems are isolated to their specific environment.
- Failures do not cross SAN island boundaries. Fault isolation is limited to each SAN.
- Firmware revisions are specific to each deployed SAN and do not have to be consistent enterprise-wide.
- Specific functions (such as mission critical applications or test environments) are isolated and do not interact with or corrupt other functions.

SAN Island Problems

Implementation and management of isolated SAN islands has several problems, including:

- A large number of fabric elements, storage devices, and servers to administer from several workstations, possibly using multiple SAN management applications.
- No economy of scale to mitigate the costs of advanced fabric features. Sophisticated management applications must be purchased to administer each SAN. If high availability and non-blocking performance are required, director-class switches must be purchased for each SAN.
- Complex interdependencies and data congestion because fabric switches are connected with multiple interswitch links (ISLs). A single low-cost edge switch can limit the scalability and performance of an entire fabric.

- Inability to consistently schedule maintenance downtime for each SAN.
- Stranded resources. Unused ports in one SAN cannot be used by applications in another (port limited) SAN, and expensive resources (such as tape backup elements) cannot be easily shared across SAN boundaries.

Large Fabric Problems

Fibre Channel fabrics configure and manage themselves, and require operator intervention only upon failure. When a Fibre Channel device connects to a director or fabric switch, the device receives a unique 24-bit (three-byte) network address composed of domain, area, and port bytes. This address is used for routing data through the fabric. Domain identifiers (Domain_IDs) are typically reserved for fabric elements (directors and switches), and range between 1 and 31 for McDATA products. A fabric element with a unique Domain_ID then allocates the remaining two bytes (area and port) to provide network addresses for devices.

The principal switch selection process ensures each fabric element has a unique Domain_ID. This process determines which director or fabric switch acts as the master in a newly-configured SAN and allocates unique Domain_IDs to the remaining elements. Without this process, two elements could have the same Domain_ID, resulting in duplicate addresses and misrouting of data.

Principal switch selection is initiated by transmitting exchange fabric parameter (EFP) frames to expansion ports (E_Ports) of all connected switches until a user-defined fabric stability timeout value (F_S_TOV) is reached. The F_S_TOV is proportional to the number of fabric elements. If the fabric is large, the F_S_TOV is set higher and the selection process takes longer.

During principal switch selection, a disruptive or non-disruptive build fabric event occurs. A non-disruptive event preserves fabric element Domain_IDs and does not require reassignment of network address blocks. However, a disruptive event may cause elements to acquire a different Domain_ID, which means each attached device must re-login to the fabric to acquire a new network address.

Ironically, the self-configuring functionality provided by the Fibre Channel architecture makes it more problematic to build large fabrics or to extend fabrics over large distances. Fibre Channel fabrics tend to become unstable long before reaching the maximum theoretical switch count (239) because of ISL congestion, disruptive build fabric events, or sporadic disruptions.

Another persistent problem associated with large Fibre Channel fabrics is multi-vendor incompatibility. Due to lack of common communications standards and fabric shortest path first (FSPF) protocol, switch vendors may have to support multiple (standards-compliant and proprietary) interoperability modes. In addition, protocol enhancements may force vendors to support multiple firmware versions for the same product. These issues make it difficult to connect directors or fabric switches from different original equipment manufacturers (OEMs) to build a large SAN.

Consolidation Solutions

McDATA offers the following solutions for SAN island consolidation:

- **FlexPar technology -** By enabling flexible partitioning (FlexPar) technology, assets can be physically consolidated while maintaining the benefits of application-based or role-based SAN islands. Refer to *Flexible Partitioning Technology* for additional information.
- **SAN routing -** By installing an Eclipse 2640 SAN Router, a collection of individual Fibre Channel fabrics is connected and then functions as a single large network (routed SAN) providing any-to-any connectivity, while maintaining the autonomous nature of individual Fibre Channel fabrics. Refer to *SAN Routing* for additional information.

Flexible Partitioning Technology

FlexPar technology enables the partitioning of fabric-attached devices to enable SAN island consolidation, decrease fabric congestion, or decrease the possibility of downtime. The technology enables partitioning through:

- **Director FlexPars -** available only for the Intrepid 10000 Director.
- **Zone FlexPars -** available for Intrepid-series directors and Sphereon-series fabric switches.
- **Role-based FlexPars -** available for Intrepid-series directors and Sphereon-series fabric switches by mid-2005.

Director FlexPars

By installing an Intrepid 10000 Director with the director FlexPar product feature enablement (PFE) key, assets can be physically consolidated while maintaining the benefits of application-based SAN islands. The feature divides the director into multiple (up to four) sub-directors, each operating with independent management and services.

This feature reduces unused ports and resources and consolidates the enterprise into a single infrastructure, while maintaining multiple independent application and fault isolation domains.

Up to four partitions can be enabled for each director (**0** though **3**), where a partition consists of one or more line modules (LIMs). User access and Fibre Channel traffic (Class 2, Class 3, and Class F) are isolated within each partition. However, the director ships with only a master FlexPar (**0**) enabled, allowing management and administration from a single point of control. Figure 4-1 illustrates director FlexPar functionality.

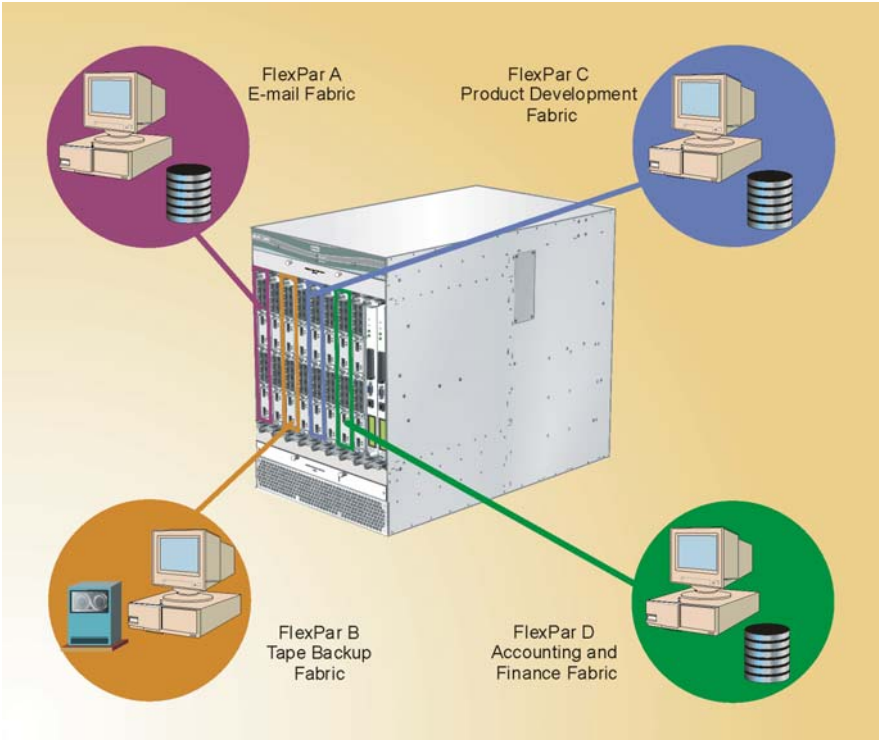

Figure 4-1 Intrepid 10000 Director FlexPar Functionality

A SAN management application or command line interface (CLI) user with read-write administrator privileges can perform the following from master Flexpar **0**:

- Create up to three additional FlexPars and assign resources to those FlexPars.
- Perform director firmware upgrades to all Flexpars.
- Enable or disable the protocol subsystem for any Flexpar.
- Enable or disable switch modules and control processor (CTP) cards.
- Perform director shutdowns, restarts, and field-replaceable unit (FRU) switchovers.
- Set the director Internet protocol (IP) address, gateway address, and subnet mask.
- Set the director date and time.

NOTE: An Intrepid 10000 Director set to fibre connection (FICON) management style can have only one Flexpar (**0**) enabled.

Flexpar functionality aggregates small (32 port or less) and medium (up to 256 port) SAN islands to a single point of control. Director FlexPars are not intended to enable creation of larger fabrics. Regardless of the number of Flexpars enabled, the maximum total fabric size (summing all Flexpars) is limited to the maximum configuration supported by the Director. Similarly, limits for other fabric scalability elements (such as number of zones per zoneset) may not exceed the maximum supported by a single fabric.

Zone FlexPars

Because zoning is managed on a fabric wide basis, all directors and fabric switches in a zone must maintain consistent zoning information. To ensure consistency, registered state change notifications (RSCNs) are transmitted to all attached devices when a zoning change occurs, when a device is set offline or online, or when a local switch is connected to or disconnected from the fabric.

When a device becomes available or unavailable, an RSCN is transmitted only to devices in the same zone. However, a zoning change causes an RSCN to be transmitted to all the devices in the fabric. As fabrics grow larger, numerous RSCNs from zoning changes can create congestion and disrupt devices, causing a pause in normal activity to determine status of the other devices. In fact, many legacy host bus adapters (HBAs) cease operation as they query the fabric name server upon receipt of an RSCN.

Zone FlexPars implement an RSCN zone isolation feature that prevents fabric-format RSCNs from propagating to devices in zones not impacted by the RSCN. With zone FlexPars enabled, zoning change RSCNs are handled like device availability change RSCNs.

Because the feature is device centric, zone FlexPars work in loop environments and with node port ID virtualization (NPIV) enabled. In addition, the feature operates when the director or switch *Interop Mode* is set to **McDATA Fabric 1.0** or **Open Fabric 1.0**.

Zone FlexPars are enabled or disabled on a per-switch basis through the CLI by setting the *zoneFlexParstate* parameter to **fabric** (enabled) or **none** (disabled). When installing a new director or switch with E/OSc 7.0 (or upgrading an existing fabric element to E/OSc 7.0 or E/OSn 6.0), the feature is enabled by default and operates on a fabric-wide basis.

Role-Based Flexpars

As Fibre Channel fabrics grow in size and complexity, the potential for fabric configuration problems caused by human error increases significantly. Implementation of role-based access control (RBAC) through role-based FlexPars (available for McDATA products by mid-2005) can control and mitigate these problems.

Through a SAN management application, users are grouped into roles, and roles are assigned a set of responsibility-based privileges. These privileges include access to specific devices and commands. In addition, roles can own subsets of a fabric. This concept is useful when a fabric includes several applications, each managed by a different administrator.

If a user is assigned administrator (role) duties for a set of switches and devices, other administrators cannot configure the user's switch and device subset. Role-based FlexPars ensure accountability for each application, fabric, or network; prevent errors from propagating across applications; and prevent unauthorized users from damaging or shutting down a fabric.

Role-based FlexPars can be configured to warn users in addition to preventing actions. Thus if one person administers an entire fabric, roles ensure the administrator is reminded of the cross-application adverse impact that a configuration action may cause.

SAN Routing

Connecting isolated, department-level, and application-specific Fibre Channel SANs is a requirement for most enterprises. Consolidating SAN islands:

- Provides campus storage connectivity and interoperability between formerly-incompatible Fibre Channel fabrics (from the same or different vendors).
- Allows construction of large Fibre Channel fabrics (up to or exceeding 239 directors or fabric switches), while reducing fabric rebuild disruptions and retaining secure partitioning of network resources through autonomous management domains.
- Provides a stable, long-distance connection over a wide area network that allows content sharing over regional distances, consolidates remote tape backup, and implements BC/DR solutions.
- Allows the enterprise to implement newer technologies and protocols (such as iSCSI) while preserving investment in a Fibre Channel infrastructure.

However, connecting Fibre Channel fabrics is not a simple process of cabling ports together. Fibre Channel architecture provides several fabric services that require attention to ensure device interoperability and stability of the consolidated SAN. A robust approach to solve this connectivity problem is secure, multi-protocol SAN routing.

A routed SAN consists of multiple Fibre Channel fabrics functioning as a single network, providing any-to-any device connectivity, but maintaining desired autonomous characteristics of individual fabrics. Storage networking devices that connect fabric elements in such a manner are called SAN Routers. Routed SANs typically include directors and fabric switches from different vendors, operating in mixed modes, using different protocols (such as Fibre Channel and iSCSI), and using several firmware versions.

Multi-protocol SAN routing provides larger-scale SAN connectivity without compromising the ease of administration, high device availability, and security inherent to SAN islands. Figure 4-2 illustrates a three-tier model that defines SAN routing hierarchy:

- **Tier 1 -** The first tier consists of isolated Fibre Channel fabrics (SAN islands). Within each fabric, data is transmitted between directors and fabric switches through E_Port ISLs.

- **Tier 2 -** To connect SAN islands without physically merging the fabrics, the second tier consists of metropolitan storage area networks (mSANs). SAN routers connect fabrics within a data center or campus to form an mSAN and transmit data between fabrics through router ports (R_Ports). Refer to *mSAN Routing* for additional information.
- **Tier 3 -** To connect geographically remote fabrics or mSANs, the third tier consists of internetworked storage area networks (iSANs). SAN routers transmit data between mSANs through intelligent ports, using Internet Fibre Channel protocol (iFCP). Refer to *iSAN Routing* for additional information.

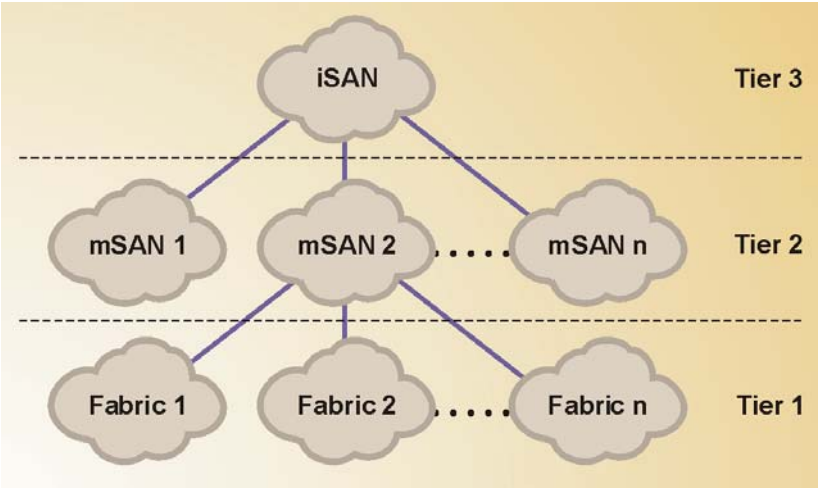

Figure 4-2 SAN Routing Hierarchy

R_Port technology enables inter-fabric SAN routing within a data center or limited geographic area to create mSANs, while iFCP inter-router links (IRLs) connect distributed, extended-distance Fibre Channel fabrics to create an iSAN. In addition, high-availability SAN routers are connected with IRLs using metropolitan Fibre Channel protocol (mFCP). Figure 4-3 illustrates these concepts.

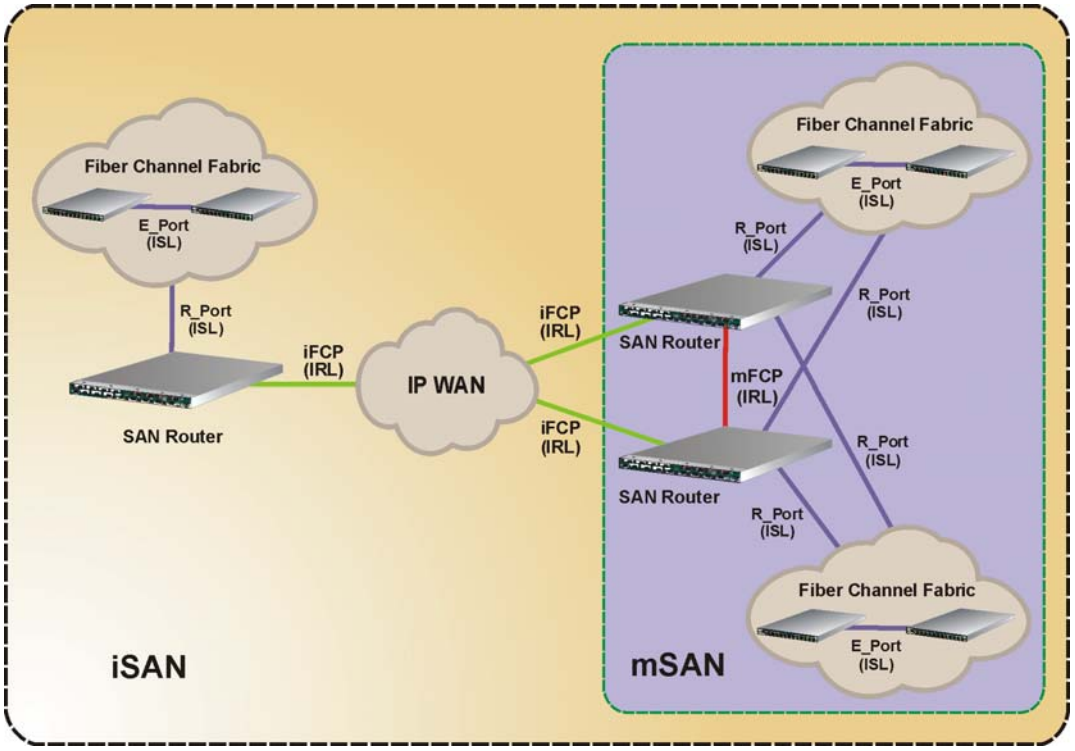

Figure 4-3 SAN Routing Concepts

The following sections discuss SAN routing concepts, including:

- R_Port operation.
- Routed SAN zoning.
- mSAN routing.
- iFCP operation.
- iSAN routing.
- Inter-FlexPar routing.
- Best practices.

R_Port Operation

To avoid building a large Fibre Channel fabric with its inherent reconfiguration issues, SAN Routing provides any-to-any connectivity (to maximize use of common assets across SAN islands), while retaining the fault isolation characteristics of smaller SANs. SAN routers also support multiple R_Port compatibility modes, making it possible to route OEM versions of a vendor switch, direct-marketed versions of a vendor switch, and switches produced by different OEMs.

An Eclipse 2640 SAN Router is used to connect multiple Fibre Channel fabrics within a data center, building, or campus. Figure 4-4 shows the example physical connectivity of Fabric **1** (one switch) and Fabric **2** (one director and one switch) through the router.

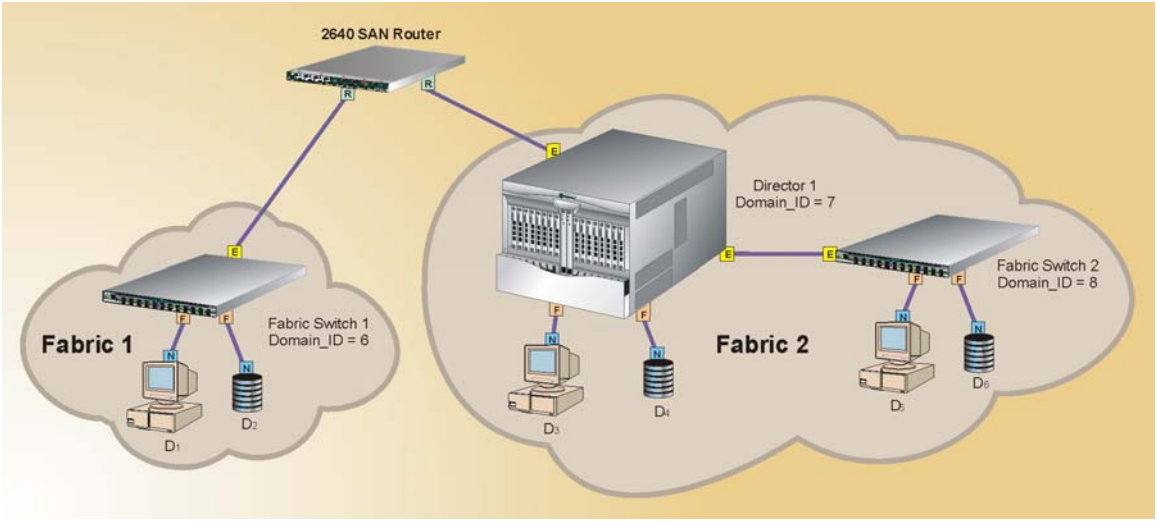

Figure 4-4 SAN Routing - Physical Connectivity

Unlike a conventional Fibre Channel E_Port, a SAN router R_Port behaves as a virtual one-port edge switch (with a unique Domain_ID) and terminates Class F traffic at the boundary of the connected fabric. Class F traffic provides control, coordination, and configuration of fabrics. Directors and fabric switches use Class F services to transmit FSPF protocol structures and related link state database information across ISLs. By terminating Class F traffic at an R_Port, switch-to-switch protocols are not passed through the router, and disruptive build fabric events are restricted to each fabric.

Instead of Class F frame transmission, routing communication is provided by Fibre Channel network address translation (FC_NAT) technology. This is similar to the technology used by IP networks to convert private addresses to public addresses.

The principal switch in each router-connected fabric assigns the Domain_ID to the associated R_Port acting as an edge switch. The switch priority for an R_Port is set to a hexadecimal value of **FF** and cannot be changed, therefore an R_Port cannot become the principal switch in the fabric.

The implication of a virtual edge switch is that each director or switch connected to a SAN router has no knowledge of other directors or switches (unless they are physically connected through E_Port ISLs). This means:

- If two Fibre Channel fabrics are connected to a SAN router, the result is not one large fabric but the two fabrics interconnected by the router. Each fabric maintains its autonomous nature.
- If multiple fabrics are routed as part of an mSAN, connecting a new fabric to the router is a nondisruptive event to the existing fabrics.
- Only authorized (zoned) connections between devices can transmit Class 2 and Class 3 Fibre Channel traffic (data frames) across routed SANs.
- Switch registered state change notifications (SW_RSCN) frames transmitted through an R_Port to the router are retransmitted to authorized (zoned) devices in another router-attached fabric.
- Each fabric has access to a full Domain_ID space, independent of other router-attached fabrics.

SAN routing provides the benefits of shared data and device access, while eliminating interoperability and fabric rebuilding issues. This enables development of complex, scalable, network storage solutions that outperform traditional Fibre Channel fabrics.

Logical Connectivity

While each R_Port is assigned a Domain_ID by the principal switch of the attached fabric, the router reserves and manages two internal routing domains with proxy Domain_IDs **30** (hexadecimal **7E**) and **31** (hexadecimal **7F**).

NOTE: Proxy Domain_IDs **30** and **31** are reserved for routing domains and cannot be assigned to directors or switches in any router-attached fabric.

The routing domain with proxy Domain_ID **30** represents Fibre Channel devices that are part of a router-attached fabric (part of a local mSAN). The routing domain with proxy Domain_ID **31** represents devices that are directly attached to the router's Fibre Channel ports or connected through an iFCP link. Figure 4-5 illustrates routing domains and shows the logical connectivity of Fabric **1** (one switch) and Fabric **2** (one director and one switch) through an Eclipse 2640 SAN Router.

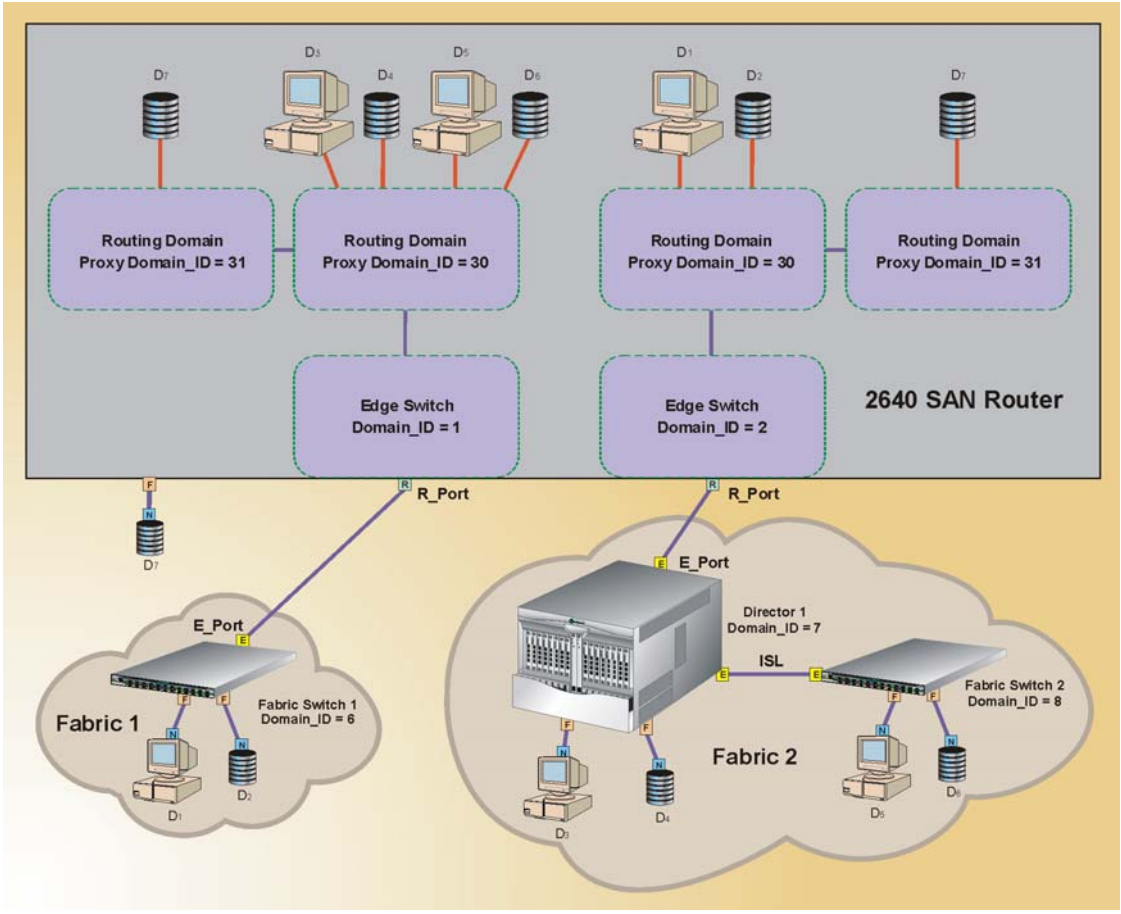

Figure 4-5 SAN Routing - Logical Connectivity

For attached fabrics in which participating element's *Interop Mode* is set to **McDATA Fabric 1.0**, Domain_IDs of **30** and **31** are recognized by SAN management applications and all attached devices. For attached fabrics in which participating element's *Interop Mode* is set to **Open Fabric 1.0**, Domain_IDs of **30** and **31** are recognized by SAN management applications. Domain_IDs of **7E** and **7F** are recognized by all attached devices.

As shown in Figure 4-5 from a logical connectivity perspective, fabric **1** appears as follows:

- One Fibre Channel switch (Domain_ID **6**) with direct-attached devices $\mathbf{D_1}$ and $\mathbf{D_2}$.
- One edge switch (Domain_ID **1**) that represents an R_Port.
- One virtual switch (Domain_ID **30**) that represents a routing domain with devices $\mathbf{D_3}$, $\mathbf{D_4}$, $\mathbf{D_5}$, and $\mathbf{D_6}$ logically attached. These devices are physically connected to Fabric **2**.
- One virtual switch (Domain_ID **31**) that represents a routing domain with device $\mathbf{D_7}$ logically attached. This device is physically connected to the router.

As shown in Figure 4-5 from a logical connectivity perspective, fabric **2** appears as follows:

- One Fibre Channel director (Domain_ID **7**) with direct-attached devices $\mathbf{D_3}$ and $\mathbf{D_4}$.
- One Fibre Channel switch (Domain_ID **8**) with direct-attached devices $\mathbf{D_5}$ and $\mathbf{D_6}$.
- One edge switch (Domain_ID **2**) that represents an R_Port.
- One virtual switch (Domain_ID **30**) that represents a routing domain with devices $\mathbf{D_1}$ and $\mathbf{D_2}$ logically attached. These devices are physically connected to Fabric **1**.
- One virtual switch (Domain_ID **31**) that represents a routing domain with device $\mathbf{D_7}$ logically attached. This device is physically connected to the router.

Routing domains represent devices from remote fabrics and are a key part of SAN routing. The routing domain with Domain_ID **30** enables routing between multiple fabrics (mSAN routing). The routing domain with Domain_ID **31** enables routing between multiple mSANs (iSAN routing). Refer to *mSAN Routing* or *iSAN Routing* for additional information.

R_Port Domain_ID Assignment

The default preferred Domain_ID for each SAN router R_Port is **1**. However, each port should be assigned a preferred Domain_ID (set at the *R_Ports* tab of the *Fabric Configuration* dialog box) that is unique within the attached fabric. The principal switch in the attached fabric then attempts to allocate this requested (preferred) Domain_ID to the R_Port. If the requested value is in use, the principal switch assigns the first available Domain_ID.

NOTE: If more than one R_Port (from the same router or multiple routers) is attached to a fabric, each port requires a unique Domain_ID.

The *Insistent Domain_ID* option (enabled at the *R_Ports* tab of the *Fabric Configuration* dialog box) ensures an R_Port gets a predictable (assigned) address. However, if the *Insistent Domain_ID* check box is enabled for an R_Port and the port does not get the preferred address, the R_Port segments and does not connect to the fabric. In addition, R_Ports segment if:

- Two R_Ports physically connect to a single fabric and the connections are configured (at the SANvergence Manager application) for attachment to a pair of fabrics.
- Two R_Ports physically connect to two fabrics and the connections are configured (at the SANvergence Manager application) for attachment to a single fabric.

It is recommended that insistent (unique) Domain_IDs be assigned to directors, fabric switches, and R_Ports in routed fabrics. Assigning Domain_IDs results in known network addresses, predictable device behavior, and ease of locating and identifying equipment.

Router Fabric Manager

A fabric-attached R_Port acts as a router fabric manager that manages fabric discovery, device registration, zoning, and other fabric-related activities between the router and attached fabric. The router fabric manager provides communication between the fabric's simple name server (SNS) and the router's metropolitan simple name server (mSNS) or Internet simple name server (iSNS). When a SAN router connects to a Fibre Channel fabric, device information is mutually exchanged. The router and the fabric's principal switch register any new device information with their respective name servers.

There is only one router fabric manager per fabric. If more than one R_Port (from the same or multiple routers) connects to a fabric, then the port with the lowest worldwide name (WWN) is elected router fabric manager for that fabric. Other R_Ports become subordinate ports. All R_Ports (router fabric manager or not) participate in FSPF routing protocol and traffic forwarding.

Routed SAN Zoning

SAN Routing provides flexibility with respect to zoning behavior and interactions between a router and attached fabrics. The zone policy (set at the *Fabrics* tab of the *Fabric Configuration* dialog box) specifies how zoning information is synchronized between the router and fabrics. It is not a requirement that all router-attached fabrics use the same zone policy. The zone policy options are:

- **No Zone Synchronization -** Device zoning is controlled at the fabric level through a Fibre Channel SAN management application. Zone configurations propagated from a SAN router to a fabric are negated through a SAN management application.
- **Append IPS Zones -** Device zoning control is shared between a SAN router (SANvergence Manager application) and a Fibre Channel SAN management application. Existing zones configured at the fabric level are synchronized through the SANvergence Manager application.

No Zone Synchronization

When the zone policy is set to **No Zone Synchronization**, zone set information between a SAN router and the associated fabrics is not synchronized. In an environment with configured Fibre Channel fabrics (prior to enabling SAN routing), it may be preferable to use the existing SAN management application to enforce zoning. This practice is applicable if all devices are fabric-attached and no devices are directly attached to SAN router ports. A typical application for this zone policy is remote data replication (RDR), where a small number of devices must communicate across fabric boundaries.

Using a SAN management application, matching zones are created in each fabric for all devices that require cross-fabric communication. This implies that two devices that need to communicate are configured in a common zone in both fabrics. At the SANvergence Manager application, one zone is created that contains all devices that are shared across the fabrics. Shared devices are visible to both fabrics through standard SW_RSCNs.

A **No Zone Synchronization** policy is typically not suitable for larger SAN routing environments where many devices must be visible to numerous fabrics.

Append IPS Zones

When the zone policy is set to **Append IPS Zones**, Internet protocol storage (IPS) zone set information from the router is appended to the active zone set for every router-attached fabric in the mSAN. In addition, all devices in the IPS zone set are added to the SNS of each fabric, even if the fabric does not have an active zone set. This policy is the default setting when an R_Port is configured.

The **Append IPS Zones** policy is recommended and provides a balance between ease of use and retention of primary zoning control by Fibre Channel SAN management applications. Cross-fabric devices are zoned together (IPS zone set) through the SANvergence Manager application, and the zone is appended to the active zone set for each fabric (zoned through a SAN management application). Any subsequent changes made to the IPS zone set are propagated to the fabrics.

NOTE: When using a SAN router with E/OSi Version 4.5 (or earlier), an active zone set must exist for each fabric prior to performing an **Append IPS Zones** operation. When using E/OSi Version 4.6 (or later), active zone sets are created automatically when the router zone is appended.

The **Append IPS Zones** policy automatically creates common router zones in each attached fabric. In large environments where many devices must be visible to numerous fabrics, this policy is much easier than creating router zones manually (using the **No Zone Synchronization** policy).

The IPS zone set appended to the active zone set of the fabric must not be modified using a Fibre Channel SAN management application. Changes made to the appended IPS zone set are immediately overwritten with the information previously configured at the SANvergence Manager application.

mSAN Routing

An mSAN consists of one or two SAN routers that interconnect up to six Fibre Channel fabrics. These fabrics are typically dispersed within a data center or metropolitan campus. If two SAN routers are used, they are connected with multiple (one to four) Gigabit Ethernet (GbE) bandwidth IRLs. These GbE connections (using mFCP as the transport protocol) are characterized by low latency, high bandwidth, and negligible packet loss.

mSAN Routing Domain

A SAN router reserves a routing domain with proxy Domain_ID **30** (hexadecimal **7E**) to enable routing between mSAN fabrics. This routing domain is visible to all directors and switches in each router-attached fabric and appears as a virtual switch (with Domain_ID **30**) to SAN management applications. Table 4-1 summarizes the mSAN routing domain.

NOTE: A reserved routing domain with proxy Domain_ID **31** (hexadecimal **7F**) enables iSAN routing.

Table 4-1 mSAN Routing Domain

Domain_ID	Area_ID	Fabric_ID
30 (Hexadecimal **7E**)	1 - 4	1
	5 - 8	2
	9 - 12	3
	13 - 16	4
	17 - 20	5
	21 - 24	6
	25 - 28	7
	29 - 32	8
	33 - 36	9
	37 - 40	10
	41 - 44	11
	45 - 48	12

During SAN router configuration, each R_Port is assigned (through the SANvergence manager application) a unique Fabric_ID between **1** and **12**. Although the theoretical limit is 12 Fabric_IDs per mSAN, the supported limit is six. As shown in Table 4-1, four Area_IDs are available to each Fabric_ID. Therefore, the combination of domain, area, and fabric IDs creates a theoretical limit of 1,024 devices per fabric (although the supported number is far less).

When a fabric element encounters a device with a Fibre Channel network address starting with Domain_ID **30** or **7E**, the associated device is physically connected to a different fabric. In addition, routing communication between the fabric element and device is provided through FC_NAT technology. Fibre Channel network addresses are not unique to each routed fabric (and require router translation for cross-fabric communication) because the Domain_ID space is reused across fabrics. Although device network addresses are router translated, device WWNs are not translated and remain consistent across the entire routed fabric.

Router Name Servers

Each SAN router in an mSAN (up to two) maintains an mSNS database. With one SAN router installed, the router maintains a primary simple name server (pSNS) database with information about all fabric-attached or router-attached devices in the mSAN (and across iSANs). The pSNS, using the router fabric manager R_Port as a conduit, interfaces with the fabric SNS to form a complete name server database.

With two SAN routers installed, one router maintains a pSNS database and the second router maintains a secondary simple name server (sSNS) database. Each mSAN always has one pSNS. The sSNS contains information only about devices directly attached to the second router and is a client to the pSNS. The pSNS router is user-selected or assigned during the build fabric process on the basis of the lowest WWN.

The secondary router sSNS transmits connectivity information to the primary router pSNS as required. Because SNS databases use unicast and subnet broadcasts to communicate, the pSNS and sSNS routers must be configured on the same subnet. If the mFCP IRL between the routers segments, different information exists in the SNS databases.

Router Connectivity through mFCP

mFCP provides connectivity (through a GbE-bandwidth IRL) between two Eclipse 2640 SAN Routers. mFCP is similar to Fibre Channel protocol (FCP) but implements user datagram protocol (UDP) for open systems interconnection (OSI) Layer 4 transport. mFCP links are used for path failover in high-availability mSANs.

NOTE: The Eclipse 1620 SAN Router does not support mFCP and must be deployed in mSANs as a single-router configuration.

The UDP transport protocol is fast and easy to implement, but unlike transmission control protocol (TCP), UDP is connectionless, best-effort, and does not guarantee order or delivery of packets. UDP does not offer services such as packet reordering, retransmission of lost packets, or detection of duplicate packets. Therefore, only direct, high-reliability fiber-optic cable connections between SAN routers are supported.

An mFCP link typically connects routers over short distances in a data center or campus. However, mFCP links can connect routers in a metropolitan area using wavelength division multiplexing (WDM) equipment or dark (unused) fiber. WDM and dark fiber are considered direct connections.

SAN routers and the UDP over GbE connection support the Institute of Electrical and Electronics Engineers (IEEE) 802.3x Ethernet flow control standard. Flow control prevents buffers from overflowing and dropping packets.

A UDP over GbE connection eliminates protocol overhead (eight bytes for UDP versus 20 bytes for TCP) and potential performance problems. The header is smaller and does not have windowing mechanisms that require resources to manage, and there is no buffering of segments until notification of receipt. The connection also uses 8B/10B bit-level encoding derived from Fibre Channel specifications, resulting in a low bit-error rate. Flow control, low overhead, and a low bit-error rate allow the mFCP connection to approach the reliability of a Fibre Channel connection.

While a SAN router IRL is limited to GbE speed, multiple IRLs can be combined using IEEE 802.3AD link aggregation standard. Up to four links can be aggregated between two SAN routers.

If a direct Fibre Channel connection exists between routed fabrics, storage traffic traverses the Fibre Channel ISL and not the router-to-router mFCP link. Only SNS traffic traverses the mFCP link. However, if a router-to-router mFCP link is the only path between two Fibre Channel devices, the link is traversed by storage traffic. Because a SAN router is the edge switch for every routed fabric and advertises Domain_ID **30** or **7E** as a direct-attach virtual switch, mFCP links do not participate in FSPF protocol in the mSAN.

mSAN Supported Limits

Table 4-2 summarizes the supported hardware and connectivity limits for a routed mSAN. Limits to the scale of an mSAN are due to inherent limits to Fibre Channel fabric SNS and SAN router pSNS databases.

Table 4-2 mSAN Supported Limits

Feature	Supported Limit
SAN routers per mSAN	An mSAN can contain up to two (**2**) Eclipse 2640 SAN Routers.
mFCP IRLs between SAN routers	Maximum number of mFCP connections between two SAN Routers is four (**4**).
Fabrics per mSAN	Maximum number of fabrics per mSAN is six (**6**). Twelve fabrics can be configured at the SANvergence Manager application, but only six are supported.
R_Port connections per fabric from all SAN routers	Maximum number of ISLs connected to a fabric from all SAN Routers (in one mSAN) is four (**4**). This provides eight Gbps bandwidth between the fabric and routers. The number of SAN routers (one or two) does not affect this limit.
Fibre Channel switches per fabric	Maximum number of directors or fabric switches per fabric is **12**. The number of SAN routers (one or two) does not affect this limit.
Fibre Channel switches per mSAN	Maximum number of directors or fabric switches per mSAN is **48**. The number of fabric elements per fabric may vary, but the total must not exceed **48**. The number of SAN routers (one or two) does not affect this limit.
Devices per fabric	Maximum number of devices per fabric is **1024**. This value is also the SNS database limit for each fabric.
Devices imported - single fabric	Maximum number of devices that can be imported from one fabric is **504**.
Devices imported - all fabrics	Maximum number of devices that can be imported to the SAN router mSNS database is **512**. All zoned devices (imported from a fabric or router-attached) are registered in the SNS databases of all connected fabrics, thus the import limit equals the fabric SNS or router mSNS database limit.

iFCP Operation

There are three protocols competing to transmit storage-related I/O traffic over long-distance transmission control protocol/Internet protocol (TCP/IP) links:

- **iSCSI** is a TCP/IP-based protocol for establishing and managing connections between IP-based storage devices, hosts, and clients. iSCSI operates on top of TCP, moving block data (iSCSI packets) over an IP Ethernet network. Refer to *iSCSI Protocol* for additional information.
- **iFCP** is a gateway-to-gateway protocol that connects distributed Fibre Channel SAN islands (or mSANs) through a TCP/IP infrastructure. With iFCP, each connected fabric is maintained separately from the others, while the IP network provides connectivity, congestion control, error detection, and error recovery.
- **FCIP -** Fibre Channel over IP (FCIP) is a TCP/IP-based protocol for connecting geographically distributed Fibre Channel SANs. FCIP requires installation of an edge device between a Fibre Channel SAN and the IP network, and encapsulates Fibre Channel frames into IP packets and fabric domains to IP addresses. This process of encapsulating one information packet inside another is called protocol tunneling. With FCIP, a single SAN fabric is created by connecting multiple SAN islands through IP network tunnels.

Typical SAN extension technologies build a single Fibre Channel fabric between two remote locations. The resulting long-distance (stretched E_Port) connection may be a direct, native Fibre Channel link (through WDM equipment or dark fiber) or an FCIP link. Using WDM equipment or repeaters, native Fibre Channel extension supports metropolitan distances up to 75 miles (120 km). FCIP supports greater distances by providing a tunneling protocol that encapsulates Fibre Channel data and forwards it over a TCP/IP network.

When two or more Fibre Channel fabrics are connected (through direct connection, WDM, or FCIP), standard fabric building and principal switch selection occurs. Whether two fabric switches are separated by a few feet or by hundreds of miles, establishing connectivity between E_Ports may trigger a disruptive or non-disruptive build fabric event. In fact, a stretched E_Port is vulnerable to disruptions caused by events at each site and to disruptions caused by problems with the extended-distance TCP/IP link.

From the standpoint of fabric build events, the only difference between a local and stretched E_Port connection is the latency introduced by the TCP/IP link and associated WDM or FCIP hardware.

A disruptive build fabric event at one local site propagates to the connected site. Likewise, a disruption in the TCP/IP link may cause the extended SAN to segment into separate SAN islands. For mission-critical storage over distance (such as disaster recovery applications) an extended SAN may inadvertently create instabilities that defeat the intent of highly-reliable data access.

iFCP operates at a higher level and addresses problems that direct connectivity and FCIP do not. iFCP is similar to FCP but uses IP for OSI Layer 3 (network layer) and TCP for OSI Layer 4 (transport layer). In contrast, mFCP uses IP for OSI Layer 3 and UDP for OSI Layer 4. Connectivity is provided through an iSNS database with a WWN-to-IP address look-up table in each fabric.

When a Fibre Channel frame is transmitted to a device in a different fabric, the frame is encapsulated and sent over the TCP/IP link to the destination fabric. The encapsulating wrapper is stripped off by iFCP and the Fibre Channel frame delivered to the destination device. iFCP can accommodate up to 64 TCP sessions per port. A TCP session opens for each pair of port WWNs that initiate a process login.

SAN routers have both FCP and IP interfaces. The Eclipse 1620 SAN router has two ports that provide IP network connectivity at up to full-duplex 100 Base-T Fast Ethernet (100 Mbps) transmission speed. The Eclipse 2640 SAN router has four ports that provide IP network connectivity at up to full-duplex GbE (1,000 Mbps) transmission speed. Figure 4-6 shows the physical connectivity of two mSANs (to form an iSAN) through Eclipse 2640 SAN Routers and an IP wide area network (WAN).

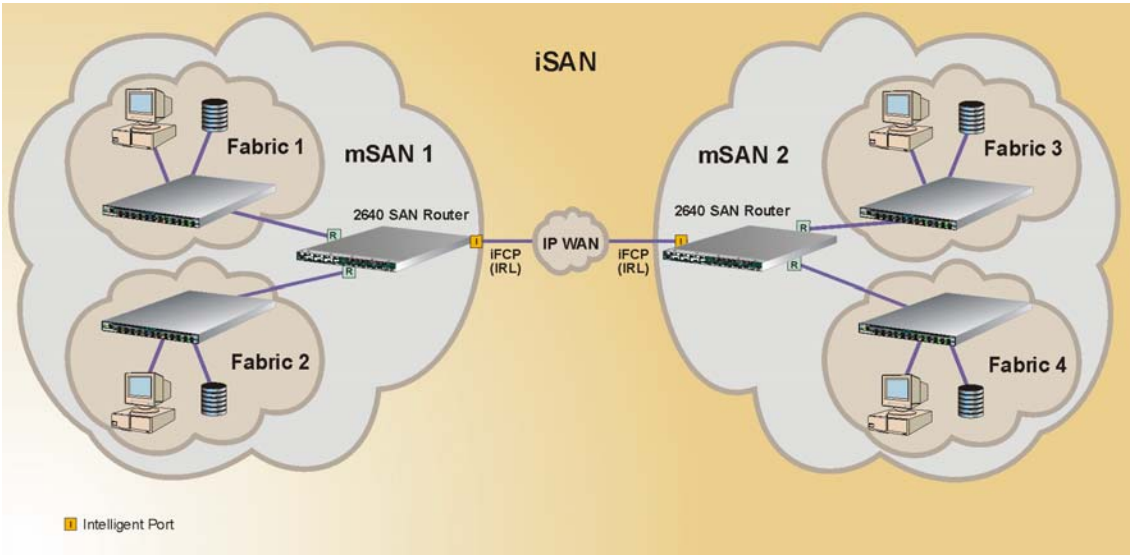

Figure 4-6 iFCP WAN Extension

iSAN Routing

An internetworked SAN (iSAN) is a network composed of multiple Fibre Channel fabrics or mSANs, connected by one or more SAN routers, where at least one fabric is remotely located and connected through a WAN. The WAN connection is an iFCP IRL, characterized by high latency and ranging in bandwidth from Digital Service 1 (DS1) at 1.544 Mbps to GbE at 1000 Mbps. iFCP is optimized for TCP / IP-based Internet service provider (ISP) networks.

Unlike conventional SAN extension, iSAN Routing terminates the stretched E_Port connection at each fabric edge. Fabric building and reconfiguration issues are isolated and restricted to each fabric because Fibre Channel Class F traffic is not transmitted across the TCP / IP network. Only authorized (zoned) connections between storage devices and servers are allowed across the network.

By preserving the autonomy of each local fabric or mSAN, an iFCP SAN routing connection ensures disruptions at one site are isolated and not allowed to propagate to other locations. This configuration provides stability for distance-connected mSANs and promotes high availability for disaster recovery, consolidated tape backup, or other complex, multi-site storage applications.

SAN Routing also streamlines SAN connectivity by eliminating network address issues associated with duplicate Domain_IDs. Because fabric elements and devices at either end of an iFCP connection remain in separate mSANs, address conflicts between the mSANs do not occur. SAN Routing provides address translation (through FC_NAT) for zoned devices with authorization to communicate across the network.

iSAN Routing Domain

The routing domain with proxy Domain_ID **31** (hexadecimal **7F**) represents devices that are directly attached to SAN router Fibre Channel ports or connected through an iFCP link. The router reserves this domain to enable routing between iSANs. This routing domain is visible to all directors and switches in each router-attached mSAN and appears as a virtual switch (with Domain_ID **31**) to SAN management applications.

NOTE: A reserved routing domain with proxy Domain_ID **30** (hexadecimal **7E**) enables mSAN routing.

When communicating with directors and switches in a specific fabric, SAN routers advertise devices associated with routing domains **30** and **31** as being directly attached to the domain, even though the devices are physically attached to a separate fabric or mSAN. Proxy Domain_ID **31** (remote mSAN over iSAN) is an internal router domain connected behind proxy Domain_ID **30** (fabric over mSAN). Therefore, if a problem occurs and there is no connectivity to routing domain **30** (hexadecimal **7E**), then there is also no connectivity to routing domain **31** (hexadecimal **7F**).

IRL Optimization

TCP is a resilient protocol that retransmits lost packets, reorders out-of-order packets, detects duplicate packets, and provides flow control mechanisms. The protocol is therefore appropriate for iFCP connectivity over extended-distance links. However, additional optimization is usually required to transport storage traffic effectively across iFCP links. These optimization methods include:

- **Rate limiting -** If ingress traffic enters the SAN router faster than egress traffic leaves, port buffers fill and cause dropped data packets. Dropped packets cause TCP to resort to internal (and inefficient) flow control, causing dramatic link throughput decrease. Rate limiting prevents this problem. Refer to *Intelligent Port Speed* for detailed information about rate limiting.
- **Data compression -** SAN router software identifies repetitive information in an output data stream and applies a compression algorithm to ensure the data is more compact and efficiently transmitted. The compression algorithm is set at the Element Manager application using the *Compression Method* drop-down list at the *Advanced TCP Configuration* dialog box. The list provides four algorithm selections:
  - **LZO -** The Lempel-Ziv-Oberhumer (LZO) compression algorithm searches for strings of characters duplicated within a block of data being compressed. Duplicated strings are removed from the data stream and replaced by an encoded string. Non-duplicate characters (literals) are output with special encoding to distinguish them from duplicate string encoding. LZO generates a self-contained compressed data block. All information needed to decompress the data is in the compressed data, and there is no history maintained by sender (for compression) or the receiver (for decompression). The algorithm is recommended when up to 64 TCP sessions are used and the available bandwidth is up to 155 Mbps (OC-3 transport level).
  - **Fast LZO with history -** This algorithm uses the LZO algorithm with a history cache. A history cache is maintained and used to more effectively compress and decompress data. The algorithm has an average compression ratio increase of approximately 20% over LZO. The algorithm is recommended when up to 8 TCP sessions are used and the available bandwidth is up to 155 Mbps (OC-3 transport level).
  - **LZO with history -** This algorithm incorporates the LZO algorithm with a history cache and Huffman encoding. Huffman encoding is an algorithm for lossless compression based on the statistical frequency of occurrence of a symbol in the file being compressed. As the probability of occurrence of a symbol increases, the compressed bit-size representation decreases. The algorithm uses additional computing resources

and results in a lower compression bandwidth. The algorithm has an average compression ratio increase of approximately 30% over LZO. The algorithm is recommended when up to 8 TCP sessions are used and the available bandwidth is between 10 Mbps (thin Ethernet) and 45 Mbps (DS3 transport level).

— **Deflate -** This algorithm incorporates a history cache with Huffman encoding. In addition, a hash table (saved in the compressed data) is used to perform string searches. Deflate is a processor-intensive algorithm with the highest compression ratio. The algorithm is restricted to use for 10 Mbps (thin Ethernet) links.

Note that a data compression ratio cannot be definitively stated, because it changes instantaneously with every data byte transmitted. A consistent byte pattern can be compressed more than a random byte pattern. For example, a defined, constant pattern can often be compressed 15:1. Already-compressed data (such as many graphic formats and some tape formats) cannot be compressed further (1:1). Most data streams are compressible from between 2:1 and 15:1, depending on the density of consistent patterns. The only way to accurately determine a compression ratio is to compress the data and measure the result.

- **FastWrite technology -** SCSI is a simplex protocol that sends a portion of a write command, then waits for a response. Multiple commands can coexist, resulting in an inefficient process on high-latency links. FastWrite is an algorithm that reduces the number of round trips required to complete a SCSI write command to one round trip. The software improves performance over WANs by mitigating the effects of latency and using the entire link bandwidth (because all data is transmitted simultaneously).

  The FastWrite algorithm responds to initiator write commands with local transfer ready (**XFR_RDY**) commands that cause the initiator to transmit an entire data set, then buffers the output data at the SAN router closest to the corresponding target device. This eliminates multiple **XFR_RDY** command transmissions and minimizes bursty data transfer over the WAN, thus reducing round-trip delays that are characteristic of extended-distance links.

mFCP to iFCP Comparison

Table 4-3 compares mFCP to iFCP and summarizes the features of each protocol.

Table 4-3 mFCP Versus iFCP

Feature	mFCP	iFCP
Purpose	LAN protocol to support short-distance SAN router connectivity	WAN Protocol to support extended-distance SAN router connectivity
OSI Layer 4 protocol	User datagram protocol	Transmission control protocol
Link latency	Low	High
Link bandwidth	High	Low
Intelligent port operation	No	Yes
FastWrite support	No	Yes
Rate limiting support	No	Yes
Data compression support	No	Yes
Provides IEEE 802.3x flow control	Yes	Yes
Provides IEEE 802.3AD link aggregation	Yes	No
SAN routers must be configured on same subnet	Yes	No

Inter-FlexPar Routing

When the FlexPar feature is enabled for an Intrepid 10000 Director, the director is divided into multiple sub-directors, each operating with independent management and services. This consolidates application-based SAN islands, but does not enable device sharing between the SAN islands (sub-directors). Inter-FlexPar routing connects a SAN router to the Intrepid 10000 Director to enable authorized (zoned) communication between FlexPars. This is a unique application of SAN routing that enables communication between sub-directors in the same physical chassis, as opposed to routing between physically separate fabric elements.

Figure 4-7 illustrates inter-FlexPar routing. Flexpar **B** (tape backup fabric) is isolated from Flexpar **C** (product development fabric) as normally desired. However, development personnel occasionally perform tape backups that require access to Flexpar **B** devices. An E_Port from each FlexPar is physically connected to a SAN router R_Port, and Flexpar **C** servers are zoned to communicate with Flexpar **B** tape devices.

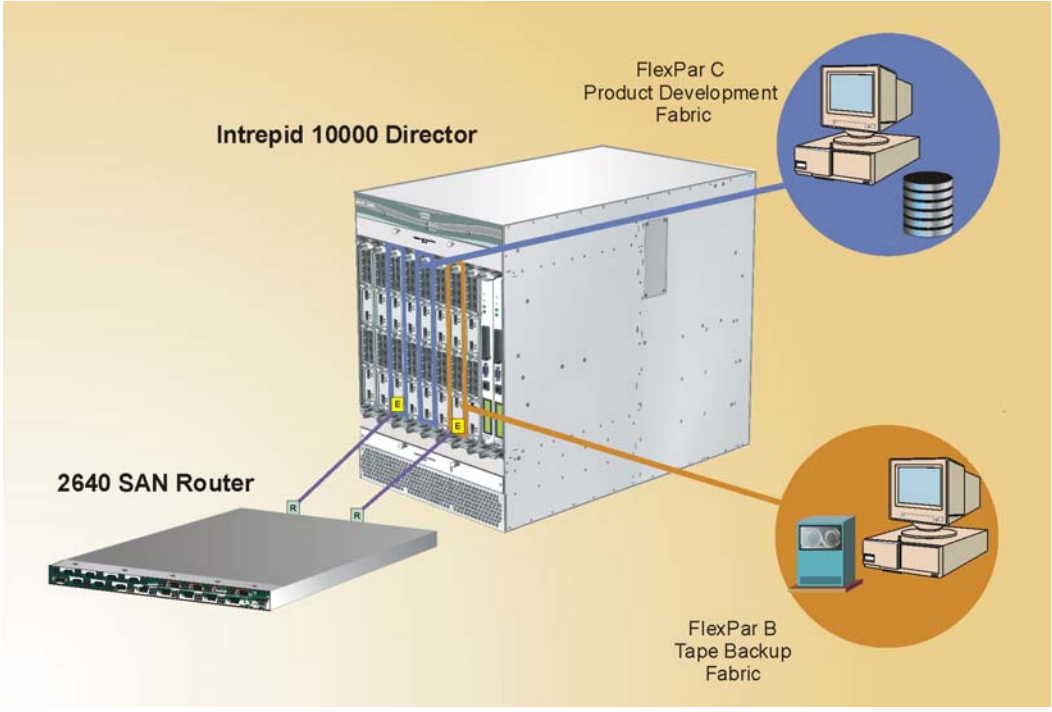

Figure 4-7 Inter-FlexPar Routing

SAN Routing Best Practices

To reduce management complexity and implement a successful SAN routing environment, follow a set of best practice conventions as follows:

1. **Plan the configuration -** Map and design the routed SAN configuration on paper, prior to installing and configuring real equipment. This includes documenting all R_Port connections, Domain_IDs, Zone_IDs, mSAN names, mSAN_IDs, iFCP port pairs, and mFCP port pairs. The connection of multiple fabrics or mSANs must be properly documented and tracked.

2. **Domain_ID assignment -** Manually assign unique Domain_IDs to all Fibre Channel directors, fabric switches, and SAN router R_Ports. Ensure the *Insistent Domain_ID* option is enabled at the SAN management or SANvergence Manager application. Do not assign Domain_ID **30** or **31** to any fabric elements. These proxy Domain_IDs are reserved for routing domains within the SAN router.

3. **Allocate Zone_IDs to mSANs -** Allocate an exclusive range of Zone_IDs for use in each mSAN. For example, each mSAN could be allocated the range **101** to **512** for local use only. These Zone_IDs are not assigned to zones shared across mSANs.

   At the *Zone Preferences* dialog box (SANvergence Manager application), leave the Zone_ID range at the default values (**1** to **512**) and track the ranges manually. If Zone_ID ranges are set, a Zone_ID outside the specified range cannot be created.

4. **Allocate Zone_IDs to shared zones -** Allocate an exclusive range of Zone_IDs for shared zones. For example, zones shared across mSANs could be allocated the range **1** to **100**. For exported zones to merge across iFCP connections, the Zone_IDs must match.

   In environments where all zones are shared zones, the local and shared Zone_ID scheme should be altered to provide different Zone_ID ranges for different mSANs. Develop a scheme where a point-to-multipoint scenario (connecting one mSAN to multiple other mSANs) carves out a range of Zone_IDs shared with each remote site. For example:

   From data center to remote site **A**: Zone_ID Range **1** to **20**.

   From data center to remote site **B**: Zone_ID Range **21** to **40**.

   From data center to remote site **C**: Zone_ID Range **41** to **60**.

5. **Assign common zone names -** Even though zone names need not be identical for zones to merge (only Zone_IDs must be identical), for simplified tracking it is good practice to assign a common name to zones that are intended to merge.

6. **Encode the Zone_ID in the zone name -** If a zone intended for database replication is assigned a Zone_ID of 7, it is good practice to include the Zone_ID in the zone name (for example DB_Replication_7). When sharing zones, this practice avoids repeatedly searching for the assigned Zone_ID.

7. **Use redundant mFCP connections -** For high availability (not increased bandwidth), use multiple mFCP connections between SAN routers to ensure the mSAN does not partition and connectivity to routing domains **30** and **31** remains intact.
8. **Assign common-numbered mFCP port pairs -** Although any FCP port can be paired with any FCP port on another SAN router, for simplified tracking it is good practice (where possible) to assign identical port numbers to both connections of a high-availability mFCP link. For example, connect port **5** of SAN router **A** to port **5** of SAN router **B**.
9. **Assign common-numbered iFCP port pairs -** Although any local intelligent port can be paired with any remote intelligent port, for simplified tracking it is good practice (where possible) to assign identical port numbers to both connections of an iFCP link. For example, connect port **14** of SAN router **A** to port **14** of SAN router **B**.
10. **Track iFCP sessions -** Every initiator-to-target device pair in a merged zone is assigned an iFCP session. Be aware of the number of active iFCP sessions. If approaching the per-port limit (64 sessions) un-export zones without active storage traffic to free up sessions.
11. **Document zones and iFCP links for each mSAN -** Use the following pair of example forms to track zone and iFCP link information. For an initial configuration, transfer values from the forms to the routed network using the SANvergence Manager Element Manager applications. For sustaining maintenance, copy information from the SAN management applications (through printable HTML reports) to the forms for consistency checks and archival.

Local mSAN Name: Boston **Local mSAN_ID: 20** **Date: 1/12/05** **Exported Zone_ID Range: 1 to 100** **Local Zone_ID Range: 101 to 512**						
Local mSAN Zone Summary						
Exported (Y/N)	**mSAN ID**	**mSAN Name**	**Description**			
Y	2	DB_Replication_2	Remote site for disaster recovery.			
Y	3	Tape_Library_3	Remote site for data center tape library access.			
N	101	Nightly_Backup	Local initiator for nightly tape backups.			
N	501	Test_Fabric	Test fabric for the local mSAN.			
iFCP Remote Connection Summary						
Local Mgmt IP Address	**Local Port Number**	**Exported Zones List**	**To Port External IP Address**	**To Port Number**	**Remote Mgmt IP Address**	**Description**
10.1.1.1	7	2	23.10.2.7	7	23.1.1.1	To Chicago.
10.1.1.1	8	3	15.2.3.7	8	15.1.1.1	To New York.

Local mSAN Name: Chicago **Local mSAN_ID: 30** **Date: 1/12/05**

Exported Zone_ID Range: 1 to 100

Local Zone_ID Range: 101 to 512

Local mSAN Zone Summary

Exported (Y/N)	mSAN ID	mSAN Name	Description
Y	2	DB_Replication_2	Replication site for disaster recovery.
N	101	Web_Server	Local web server and associated hardware.
N	501	Local_Test	Local test fabric.

iFCP Remote Connection Summary

Local Mgmt IP Address	Local Port Number	Exported Zones List	To Port External IP Address	To Port Number	Remote Mgmt IP Address	Description
23.1.1.1	7	2	10.1.1.7	7	10.1.1.1	To Boston.

12. **Configure R_Ports -** For all configured SAN router R_Ports in the same fabric:

    a. The R_Port interconnect modes (**McDATA Fabric 1.0** or **Open Fabric 1.0**) must be identical. This parameter is set at the SANvergence Manager application. The corresponding E_Port interoperability mode must also be identical. This parameter is set at the director or fabric switch Element Manager application.

    b. The zone policies (**No Zone Synchronization** or **Append IPS Zones**) must be identical. This parameter is set at the SANvergence Manager application.

    c. The error detect time-out values (ED_TOVs) must be identical. This parameter is set at the director or fabric switch Element Manager application.

    d. The resource allocation time-out values (RA_TOVs) must be identical. This parameter is set at the director or fabric switch Element Manager application.

    e. The port Domain_IDs must be different. This parameter is set at the SANvergence Manager application.

13. **Multi-vendor guidelines -** SAN routers support existing multi-vendor fabrics. However, when building a new fabric, it is good practice not to mix director and fabric switch vendors within the same fabric. A homogeneous fabric simplifies operation and fault isolation.

14. **Port zoning -** Comply with the following rules and best practices when using port zoning in a SAN routing environment:

    a. When a Fibre Channel fabric connects to a SAN router, port zoning is implemented per FC-SW2 standard. When communication between devices in a fabric stays within the fabric, port zoning works normally.

    b. When a router zone policy is set to **No Zone Synchronization**, remote devices are not port zoned in a local fabric. This is because remote entries in the local mSNS are node (device) WWNs, not port WWNs. For example, when a port-zoned device in Fabric A is zoned with a device in Fabric B, the device from Fabric A is WWN zoned. Although router ports can be port zoned, devices connected to a router port are propagated to remote fabrics using node (device) WWNs.

c. Port zoning can be confusing in a multi-vendor environment because OEMs implement zoning in different ways. When zoning with vendor-specific SAN management applications, zone through port WWNs, not node (device) WWNs.

d. Ensure devices are physically connected before importing their node (device) WWNs to a router. If a fabric port is port-zoned with nothing connected to the port, the zone member is invisible to the router until a device is connected to the port and explicitly imported to the router. If a device is disconnected or reconnected to a port-zoned fabric port, a zone update is not generated at a remote fabric (to remove the associated WWN-zoned device).

e. If router-attached directors and fabric switches have zoning licences and the zone policy is set to **Append IPS Zones** at the SANvergence Manager application, all zone licences must be enabled.

f. If router-attached directors and fabric switches have zoning licences and the zone policy is set to **No Zone Synchronization** at the SANvergence Manager application, some fabric elements may be able to operate with the zone licences disabled. Operation is vendor-specific.

g. When a device is zoned through the router CLI or SANvergence Manager application, the device is visible to all router-attached fabrics. When a device is zoned through a Fibre Channel SAN management application, the device is invisible to remote fabrics.

h. Some vendor-specific SAN management applications cannot display devices outside the local fabric. In such a case, the zone policy must be set to **Append IPS Zones** at the SANvergence Manager application.

15. **Feature conflicts -** SAN routers cannot attach to McDATA fabrics with the SANtegrity Binding feature (including both fabric binding and switch binding), OpenTrunking feature, or *Enterprise Fabric Mode* enabled. These features must be disabled before connecting the router. In addition, SAN routers do not support FICON cascading or FICON routing.

Implementing BC/DR Solutions

The post-9/11 business environment requires corporations to protect critical data by implementing cost-effective business continuity and disaster recovery (BC/DR) solutions. These BC/DR solutions drive the requirement to extend local data center SANs to geographically distant locations.

The business case for SAN distance extension is the high cost of downtime, a period during which a corporation cannot generate revenue due to temporary (or permanent) loss of critical applications or data. By connecting SAN islands through an extended-distance optical network, the corporation:

- Preserves valuable information assets and protects against business disruptions caused by facility outages, IT or communication problems, natural disasters, or terrorism.
- Provides real-time disaster recovery of business data and the IT infrastructure in the event of an unplanned outage.
- Consolidates storage resources, increases the availability of critical information, and reduces backup and restore times.
- Complies with regulatory, data protection, and data retention requirements imposed by the government and business insurers.

BC/DR solutions impose distance extension requirements to connect SAN islands. Extended-distance data transmission imposes different communication and protocol requirements. Differences between storage traffic through a local SAN and network traffic through an extended-distance WAN include:

- **Protocol stack -** Software protocol stacks quickly overload servers and inhibit SAN performance. Therefore, SANs are usually based on FCP optimized for storage environments, offering high-speed and low-overhead communication. Data networks rely on a protocol stack to provide communication and are often implemented using TCP/IP over GbE. TCP/IP provides a high level of protocol processing and is appropriate for data networks.
- **Latency -** Local storage traffic requires minimal delay and latency. Distance transmission associated with WANs introduces variable delays and high latency.

- **Reliability -** Local storage traffic requires high-reliability communication and is intolerant of data loss, out-of-order packet receipt, or data retransmission. WANs typically provide best-effort communication service and rely on upper-level protocols for end-to-end transport.

Because of these differences, a protocol conversion approach is usually required to integrate Fibre Channel SAN traffic over a geographically-dispersed network. Refer to *SAN Extension Transport Technologies* for detailed information about native FCP distance extension or protocol conversion.

Other BC / DR requirements vary, depending on budget, data type, data volume, business situation, and SAN applications. Analyze and understand these business requirements prior to selecting an operational mode (described in *Extended-Distance Operational Modes*) and transport technology (described in *SAN Extension Transport Technologies*) that best support the SAN distance-extension strategy. In particular, consider:

- **Data priority -** Not all data is critical to immediate business resumption. Prioritize and categorize data as mission-critical, secondary, or only to be retained for legal purposes.
- **Recovery time objective (RTO) -** The RTO is the time required to restore the data and applications following an outage or disaster. The loss of revenue from suspended business operations drives this objective.
- **Recovery point objective (RPO) -** The RPO is the time between backup points and defines how far out of date a backup copy can be after a failure. The data rate of change and cost of destroyed data (between the last backup and a disaster) drive this objective.
- **Distance -** The distance between a data center and replication site is proportional to RTO and RPO times. As distance between sites increases, time required to perform backup and restore operations also increases.

Extended-Distance Operational Modes

A primary component of business continuance and disaster recovery is remote data replication to an alternative safe location. Extended-distance operational modes are:

- **Synchronous remote data replication (RDR/S) -** This operational mode ensures a remote data copy (identical to the primary copy) is created at the time the primary data is created. An update operation does not complete until confirmed at both the primary and mirrored sites. An incomplete operation rolls back at both locations, ensuring the remote copy is a mirror image of the primary copy. RDR/S is synonymous with disk mirroring.

  The advantage to using RDR/S is quick data recovery. Operation at a remote, mirrored site begins immediately if operation at the primary site is disrupted. The problem with RDR/S is distance limitation. Although propagation of laser light pulses can theoretically extend to infinity, latency is an issue because propagation delays lengthen with increased link distance. These delays adversely impact performance by forcing an application to wait for confirmation of I/O operation at local and remote sites. This means RDR/S operation is distance-limited, depending on application response time tolerance and other factors.

  RDR/S is ideal for shorter metropolitan distances and real-time disk mirroring and is an appropriate BC/DR solution for enterprises requiring fast data recovery, minimal data loss, and protection against database integrity problems.

- **Asynchronous remote data replication (RDR/A) -** This operational mode does not require a response indicating completion of a remote transaction before local I/O operations resume. Replication software on the remote storage array controller ensures data is successfully written to the remote site. Standard disk backup and tape vaulting are RDR/A operations.

  RDR/A may lose data transactions during an unplanned failover to a remote site. However, after post-outage transaction logs are applied to the remote data image, operations can usually resume.

  Catastrophic events can occur anywhere. However, it is unlikely an event will span a large geographical area or two exclusive events will simultaneously occur in two locations. Therefore, to protect critical data (required for business continuance), it is prudent to replicate the data at a remote site thousands of miles from the primary site. Because there is no propagation delay involved in confirming remote transactions, RDR/A can span virtually any geographical distance and is ideal for long-distance BC/DR applications.

SAN Extension Transport Technologies

There are several extension transport technologies available to connect geographically-dispersed SAN islands, all of which differ in performance, latency, and implementation cost. The primary technologies include:

- Dark fiber (repeated or unrepeated).
- Wavelength division multiplexing (WDM).
- Synchronous optical network (SONET) and synchronous digital hierarchy (SDH).
- Internet protocol (IP).

Dark Fiber

Dark fiber refers to an installed fiber-optic infrastructure (including cabling and possibly including repeaters) that is not in use. Dark fiber strands (usually deployed in transmit and receive pairs) provide point-to-point, unprotected connectivity between two locations.

Many corporations install excess fiber-optic cabling with the expectation of leasing the infrastructure at a future date. When a telecommunication company installs cable, they often lay additional (unused) cables to avoid retrenching costs. Utility companies often install unused cables coincident with pipelines or electrical power lines. The dark fiber is then leased to companies (dark fiber service) that require dedicated optical connectivity between separate locations. Cable operation and connectivity are not controlled by the service provider. The service lessee is responsible for laser transceivers and other equipment that make the cabling functional.

Figure 4-8 illustrates extended-distance connectivity through a dark fiber (dedicated FCP or FICON) interface. The technology:

- Is well suited as an extension technology for RDR/S applications over metropolitan distances up to 22 miles (35 km) without repeaters and up to 75 miles (120 km) with repeaters. The supported bandwidth is dependent on fiber-optic quality and the choice of multiplexing scheme.
- Requires sufficient buffer-to-buffer credits (BB_Credits) assigned to the link (such as credits available through the Intrepid 10000 Director buffer pool). Refer to *Distance Extension Through BB_Credit* for information about requirements.

- Creates one logical Fibre Channel fabric through a stretched E_Port connection. The connection is vulnerable to disruptions caused by events at each site or to disruptions caused by problems with the extended-distance dark fiber link.

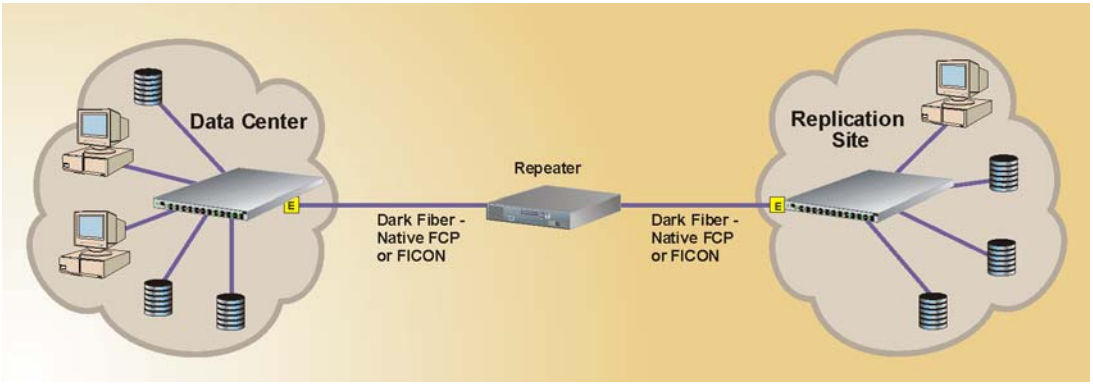

Figure 4-8 Dark Fiber Extended-Distance Connectivity

Due to the high cost of burying cables, dark fiber has limited physical availability relative to other WAN link options. Because dark fiber is usually buried, it is often susceptible to damage from excavating equipment. For these reasons, dark fiber is used on a limited basis for metropolitan-distance RDR/S applications. Dark fiber is not practical for RDR/A applications that span large distances.

WDM

Optical networks consist of fibers transmitting laser-generated flashes of light, and more information is transmitted by increasing the number of flashes per second (increasing the bit-rate). Using multiple lasers to simultaneously transmit different colors (wavelengths) of light also increases the capacity of optical fibers. Assigning laser light to designated frequencies, multiplexing (combining) the result to one signal, and transmitting the signal one fiber is called wavelength division multiplexing. At the receiving end, combined wavelengths are separated (demultiplexed). Each wavelength requires a discrete detector to convert light pulses to useful information.

The number of wavelengths used is a power of two (2, 4, 16, 32, 64, or eventually more). Technology that provides 64 wavelengths or more per fiber is called dense wavelength division multiplexing (DWDM). Technology that provides 32 or fewer wavelengths per fiber is called coarse wavelength division multiplexing (CWDM). CWDM is less complex and expensive to deploy.

Light wavelengths used are typically around 1,550 nanometers (nm). Optical fiber performs well in this wavelength region, with very little attenuation. For CWDM, differing wavelengths are separated by multiples of 20.0 nm. For DWDM, differing wavelengths are separated by multiples of 0.8 nm. The lower wavelength numbers provided by CWDM are due to lower accuracy (and price) of lasers. DWDM wavelengths are spaced closer together and require more precise lasers to reduce interference between wavelengths.

CWDM and DWDM are metropolitan extension technologies that transmit data parallel-by-bit or serial-by-character over a fiber-optic network. The signal is never terminated in the optical layer and is therefore bit-rate and format independent. As a result, WDM provides high bandwidth, low latency, and transparency to SAN protocols and allows transmission of e-mail, voice, video, multimedia, and digital data over native FCP or FICON links. Figure 4-9 illustrates extended-distance connectivity through a WDM interface.

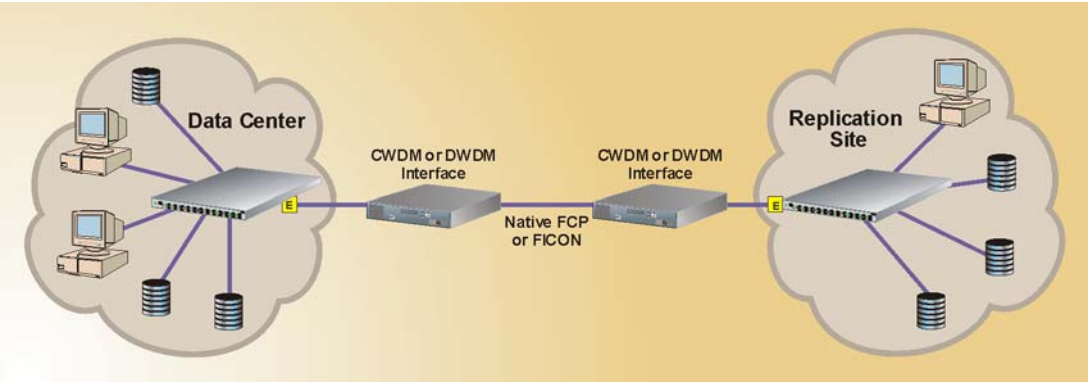

Figure 4-9 WDM Extended-Distance Connectivity

When combined with a dedicated FCP or FICON link, the technology:

- Is well suited as an extension technology for RDR/S applications over metropolitan distances up to 75 miles (120 km). Note that WDM technology does not increase the transmission distance provided by repeated dark fiber. However, WDM significantly increases the bandwidth.

- Requires sufficient BB_Credits assigned to the link (such as credits available through the Intrepid 10000 Director buffer pool). Because WDM is a method to transmit multiple signals over the same fiber-optic cable, there is no BB_Credit limitation difference between WDM and dark fiber. Refer to *Distance Extension Through BB_Credit* for information about requirements.
- Creates one logical Fibre Channel fabric through a stretched E_Port connection. The connection is vulnerable to disruptions caused by events at each site or to disruptions caused by problems with the extended-distance WDM link.

Several network service providers provide metropolitan and long-distance (intercity) WDM transport services. WDM service can be purchased on a monthly basis in accordance with a negotiated service level agreement (SLA). A typical SLA specifies the availability, minimum dedicated bandwidth (usually scalable), latency, security level, monitoring level, packet loss, and mean time to repair (MTTR).

Like dark fiber, WDM is not practical for long-distance RDR/A applications. In addition, WDM technology is still evolving and equipment is relatively expensive.

SONET and SDH

The telecommunications industry developed SONET and SDH standards for transport of time division multiplexed (TDM) data over fiber-optic cable. SONET is used in North America (United States and Canada) and Japan. SDH is used elsewhere, primarily in Europe. SONET and SDH are closely related standards that specify interface parameters, rates, framing formats, multiplexing methods, and management for synchronous TDM data transport.

The physical-layer protocol multiplexes ***n*** incoming bit streams, optically modulates the result, and transmits the signal at a rate equal to the incoming bit rate times ***n***. As an example, four data streams arriving at a SONET multiplexer at 2.5 Gbps are transmitted as one stream at 10 Gbps (four times 2.5 Gbps). Low bit-rate streams of information are multiplexed into higher bit-rate streams and transmitted at the rate of the SONET or SDH network. TDM ensures a constant stream of data through a network and takes advantage of the available bandwidth.

SONET and SDH are globally standardized technologies, more widely deployed than dark fiber or WDM, and provide a protected connection between two locations. SONET and SDH rings are also self-healing. This means a link is usually restored within 50 ms of break detection without user intervention. This makes SONET and SDH highly-available services.

Generic Framing Procedure (GFP) is a protocol-independent SONET and SDH standard that defines a mapping scheme for storage protocols such as native FCP, FICON, and iFCP. The standard includes a forward error-correction scheme that enables low bit-error rates critical for storage connectivity. Additionally, the protocol mapping to SONET or SDH requires little overhead and has minimal impact on latency and throughput. GFP is deployed with virtual concatenation (VCAT), a standard that increases bandwidth efficiency by flexibly extending bandwidth allocation in 50-Mbps increments. This extends the bandwidth range from 50 Mbps to full Fibre Channel rates. To support storage extension over long distances, GFP provides buffering and flow control to ensure high throughput without the need for Fibre Channel BB_Credit buffering. Figure 4-10 illustrates multiple extended-distance connections through a SONET interface.

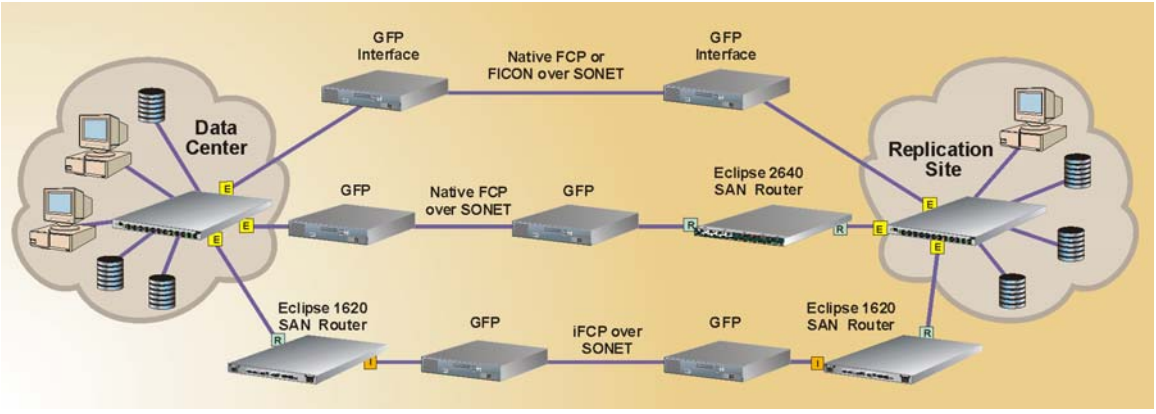

Figure 4-10 SONET Extended-Distance Connectivity

The technology:

- Is widely deployed and highly-available, and is well suited as an extension technology for RDR / A applications over thousands of kilometers.

- Does not require Fibre Channel BB_Credits assigned to the link because buffering is built in to the GFP function to enable long-distance transmission of storage traffic.
- Creates a routed iSAN or one logical Fibre Channel fabric through a stretched E_Port connection, depending upon the protocol and if one or more SAN routers are deployed in the link.
  - **Native FCP or FICON (unrouted) -** The top data path in Figure 4-10 illustrates native FCP or FICON extended-distance connectivity through an unrouted SONET interface. GFP interfaces provide required link buffering and flow control. The stretched E_Port connection is vulnerable to disruptions caused by events at each site or to disruptions caused by problems with the extended-distance SONET link.
  - **Native FCP (routed) -** The middle data path in Figure 4-10 illustrates native FCP extended-distance connectivity through a routed SONET interface. GFP interfaces provide required link buffering and flow control. A single Eclipse 2640 SAN router at one end of the link provides R_Port connectivity and SAN isolation (not intelligent port connectivity or protocol conversion). The routed SAN connection ensures disruptions at one site are isolated and not allowed to propagate to other locations. This connection does not support FICON operation.
  - **iFCP (routed) -** The bottom data path in Figure 4-10 illustrates iFCP extended-distance connectivity through a routed SONET interface. GFP interfaces provide SONET connectivity but are not required for link buffering and flow control. Eclipse 1620 SAN routers at each end of the link provide intelligent port connectivity and iFCP protocol conversion. The routed iSAN connection ensures disruptions at one site are isolated and not allowed to propagate to other locations. This connection does not support native FCP or FICON operation.

Several network service providers provide metropolitan and intercity SONET and SDH transport services. Long-distance SONET and SDH circuits are common and the technology does not suffer the cost and availability restrictions inherent to dark fiber and WDM. The technology provides low overhead, high bandwidth, point-to-point transport of storage traffic, and is a cost-effective choice for remote data replication.

SONET or SDH service can be purchased on a monthly basis in accordance with a negotiated SLA. However, the transport links may require sufficient BB _Credits to use the purchased bandwidth. Because of BB_Credit limitations, GFP equipment must provide buffering and flow control for native FCP or FICON storage data. Without GFP equipment, iFCP must be used to transmit the data.

Internet Protocol

As demonstrated by the Internet, an infrastructure based on IP and Ethernet delivers an unrestricted topology that scales to large geographical distances. IP and Ethernet have well-developed cross-vendor capabilities, routing, and security, and there is no inherent distance limitation. This means storage over IP (SoIP) is well-suited to provide the low to medium bandwidth (over longer distances) required for RDR / A. Block-based SoIP protocols include:

- **iFCP -** This protocol uses FCP to provide SCSI command set encapsulation, enabling Fibre Channel device communication across an IP network.
- **iSCSI -** This protocol encapsulates the SCSI command set directly to the IP transport network without relying on Fibre Channel conventions. Refer to *iSCSI Protocol* for additional information.

iFCP is an application-layer gateway protocol (FCP-to-iFCP-to-FCP) solution that connects remote storage devices or SANs across extended distances that Fibre Channel cannot support. iFCP effectively replaces a Fibre Channel SAN with an IP network but continues storage application support.

The protocol can be used over the Internet or a dedicated GbE network (with IP traffic engineering). Each connected Fibre Channel fabric (or SAN) is maintained separately, while the IP or GbE network provides connectivity, congestion control, error detection, and error recovery. Figure 4-11 illustrates SoIP extended-distance connectivity. The technology:

- Is widely deployed and highly-available, and is well suited as an extension technology for RDR / A applications over thousands of kilometers.
- Does not require Fibre Channel BB_Credits assigned to the link. IP and GbE transmit data frames without use of BB_Credits and enable long-distance transmission of storage traffic.
- Creates a routed iSAN through an extended-distance iFCP interface. Eclipse 1620 SAN routers at each end of the link provide intelligent port connectivity and iFCP protocol conversion. The

routed SAN connection ensures disruptions at one site are isolated and not allowed to propagate to other locations. This connection does not support native FCP or FICON operation.

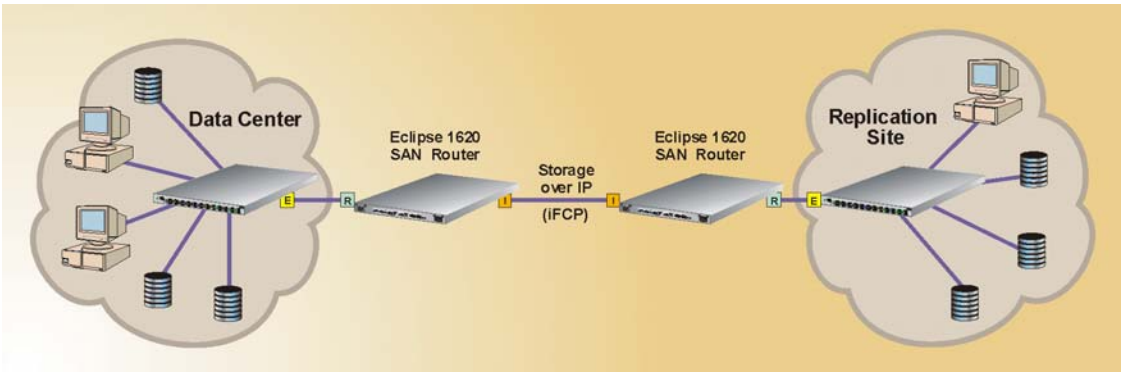

Figure 4-11 SoIP Extended-Distance Connectivity

Several network service providers provide long-distance IP or GbE network transport services. Long-distance circuits are common. The technology provides low overhead, low to medium bandwidth, point-to-point transport of storage traffic, and is a cost-effective choice for remote data replication.

Technology Comparison

Figure 4-12 illustrates the complex relationship between RTO, RPO, and extended-distance transport technology options. Extended-distance operational modes (RDR/S and RDR/A) are directly associated with the transport technology choice. Note there is substantial overlap in the functionality provided by transport technologies and no single transport technology satisfies all BC/DR requirements. Comparison factors to consider are:

- **Repeated or unrepeated dark fiber -** This technology supports medium-bandwidth, low-latency applications with short RTO and RPO requirements. Applications include real-time disk mirroring (RDR/S or RDR/A) over short to medium metropolitan distances. Unless one or more SAN routers are included in the extended-distance link (native FCP only), the technology is vulnerable to disruptions caused by fabric or link problems.

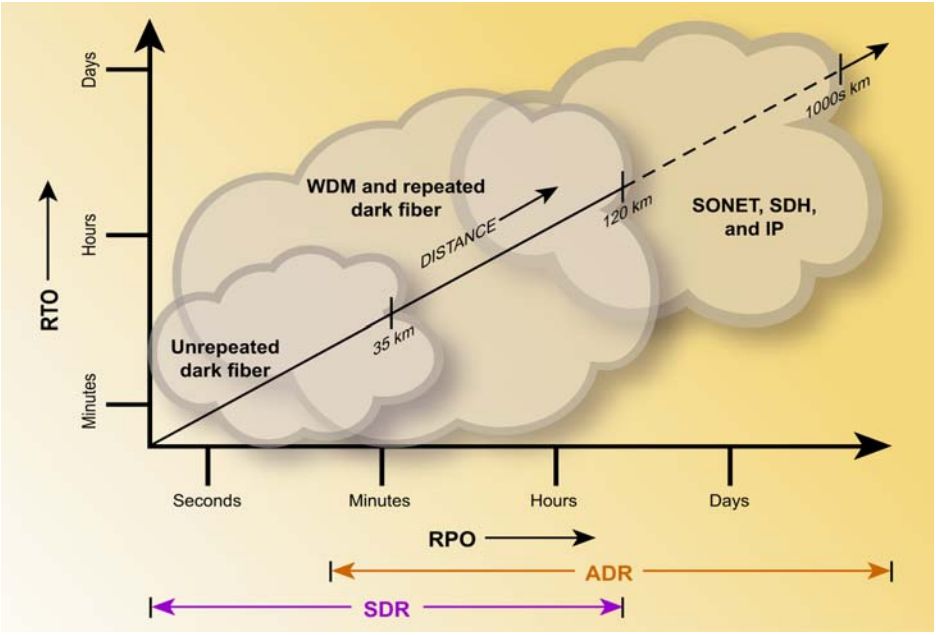

Figure 4-12 SAN Extension Technology Comparison

- **WDM -** This technology supports high-bandwidth, low-latency applications with short RTO and RPO requirements. Applications include peer-to-peer computer clustering (grid computing) and real-time disk mirroring (RDR/S or RDR/A operational mode) over short to medium metropolitan distances. WDM scales to higher bandwidths at a lower relative cost than SONET or SDH. Unless one or more SAN routers are included in the extended-distance link (native FCP only), the technology is vulnerable to disruptions caused by fabric or link problems.
- **SONET and SDH -** These technologies support medium-bandwidth, medium-latency applications with short-to-long RTO and RPO requirements. Applications include asynchronous disk backup or tape vaulting over metropolitan to extended (intercity) distances. Unless one or more SAN routers are included in the extended-distance link (native FCP or iFCP), the technology is vulnerable to disruptions caused by fabric or link problems.

- **IP -** This technology supports low-bandwidth, high-latency applications with long RTO and RPO requirements. Applications include asynchronous disk backup or tape vaulting over metropolitan to extended (intercity) distances. SAN routers are included in the extended-distance link (iFCP only), so the technology isolates the connected SANS and prevents disruptions caused by fabric or link problems.

Table 4-4 compares and contrasts the transport technologies.

Table 4-4 Transport Technology Comparison

Requirement	Dark Fiber	WDM	SONET/SDH	IP
Bandwidth (native storage)	Medium	High	Medium	Low
Extended-link latency	Low	Low	Medium	High
Network scalability	Fair	Good	Fair	Good
Performance monitoring	Average	Average	Good	Good
Extended distance (greater that 120 Km)	No	No	Yes	Yes
Security	Good	Good	Good	Good
Routed SAN benefits	No	No	No	Yes

Additional factors to consider are:

- **Availability of a physical infrastructure -** If fiber-optic cable is available, WDM is a good choice because of high bandwidth, low cost, and ease of use. SONET and SDH connectivity is generally available within metropolitan and intercity regions. IP provides the highest level of long-distance connectivity but supports only low-bandwidth, high-latency applications.
- **Bytes of data requiring backup -** The volume of data associated with the SAN is a consideration in selecting the transport bandwidth. As an example, the approximate time required to perform a 60-terabyte backup is:
  - 50 days over a single OC-3 connection.
  - One week over a single GbE connection.
  - Three hours over a 2 Gbps, 32-channel WDM connection.

- **SAN routing requirements -** If a single logical Fibre Channel fabric (created through a stretched E_Port connection) is unacceptable because of the potential for disruptive fabric rebuilds, include one or more SAN routers in the extended-distance link.

Distance Extension Through BB_Credit

Token-based buffer-to-buffer flow control governs transmission of data and link control frames in a Fibre Channel switched fabric. To manage flow control, Fiber channel fabric ports (F_Ports) or node ports (N_Ports) for the Intrepid 10000 Director are assigned a variable number of BB_Credits. The credits are typically user-defined or set to a default value. A frame cannot be transmitted through a fabric unless accounted for by a buffer credit.

During fabric port login, two communicating Fibre Channel ports exchange available BB_Credit information. When a data or link control frame is transmitted, the BB_Credit count for the port is decremented. When the transmitting port receives a corresponding receive-ready (R_RDY) link control frame, the BB_Credit count for the port is incremented.

Longwave laser transceivers and a sufficient allocation of BB_Credits are required to support long-distance transmission of Fibre Channel data frames (up to 35 km). Installation of repeaters or DWDM equipment is required to support data transmission in excess of 35 km. In either case (repeated or unrepeated) sufficient BB_Credits are required.

BB_Credits required to support a specific transmission distance are also a function of data rate. The faster the data rate, the shorter the distance supported. Approximate relationships between data rate, BB_Credits, and transmission distance are:

- At 1.0625 gigabits per second (Gbps), one BB_Credit supports a two km transmission distance (1:2 ratio).
- At 2.1250 Gbps, one BB_Credit supports a one km transmission distance (1:1 ratio).
- At 10.2000 Gbps, six BB_Credits support a one km transmission distance (6:1 ratio).

To support greater Fibre Channel transmission rates (long-link ports), the Intrepid 10000 Director provides a buffer pool that allocates user-defined BB_Credits to each port. This buffer pool is increased if the remote fabric PFE key is enabled (refer to *Remote Fabric* for information). Each director line module (LIM) contains two scalable packet processors, each supporting two optical paddles. A paddle pair provides 16 ports (1.0625 or 2.1250 Gbps operation), four ports (10.2000 Gbps operation), or ten ports (mixed data rate operation). The buffer pool is allocated among all ports in a paddle pair subject to the following constraints:

- Each paddle pair is allocated a maximum of **1,373** BB_Credits.
- Each short-link 1.0625 or 2.1250 Gbps port must be assigned a minimum of **16** BB_Credits (default value).
- Each short-link 10.2000 Gbps port must be assigned a minimum of **96** BB_Credits (default value).

Users assign BB_Credits to ports at the Element Manager application, using entries in the *RX BB Credit* column of the *Configure Ports* dialog box. Assuming four port paddle-pair combinations, the following examples explain BB_Credit allocation to configure one port in a paddle pair for extended distance (long-link) operation:

- **1.0625 or 2.1250 Gbps long link -** A paddle pair with 1.0625 or 2.1250 Gbps ports provides 16 connections. **16** BB_Credits are assigned to 15 short-link ports (**240** BB_Credits total). The remaining **1133** BB_Credits are assigned to the long-link port, supporting a repeated transmission distance of 2,200 km (1.0625 Gbps) or 1,100 km (2.1250 Gbps).
- **10.2000 Gbps long link -** A paddle pair with 10.2000 Gbps ports provides four connections. **96** BB_Credits are assigned to three short-link ports (**288** BB_Credits total). The remaining **1085** BB_Credits are assigned to the long-link port, supporting a repeated transmission distance of 180 km.
- **1.0625 or 2.1250 Gbps long link (mixed paddles) -** A paddle pair with one 1.0625 or 2.1250 Gbps paddle and one 10.2000 Gbps paddle provides ten connections. **16** BB_Credits are assigned to the slowest seven short-link ports, and **96** BB_Credits are assigned to both 10.2000 Gbps short-link ports (**304** BB_Credits total). The remaining **1069** BB_Credits are assigned to the long-link port, supporting a repeated transmission distance of 2,000 km (1.0625 Gbps) or 1,000 km (2.1250 Gbps).

- **10.2000 Gbps long link (mixed paddles) -** A paddle pair with one 1.0625 or 2.1250 Gbps paddle and one 10.2000 Gbps paddle provides ten connections. **16** BB_Credits are assigned to the slowest eight short-link ports, and **96** BB_Credits are assigned to one 10.2000 Gbps short-link port (**224** BB_Credits total). The remaining **1149** BB_Credits are assigned to the long-link port, supporting a repeated transmission distance of 190 km.

Intelligent Port Speed

When data ingress exceeds data egress for a network device, the device buffers fill, overflow, and drop data packets. Dropped packets are retransmitted. However, TCP flow control interprets these dropped packets as congestion and closes the TCP segment window. Cyclically closing and re-opening segment windows is inefficient and results in dramatically reduced link throughput. Figure 4-13 illustrates this phenomenon (data ingress exceeds egress).

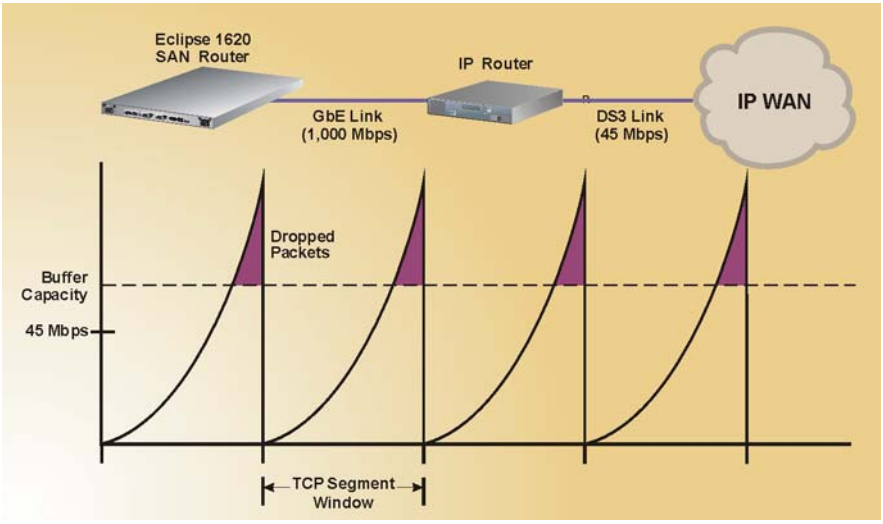

Figure 4-13 WAN Link Performance (No Rate Limiting)

To prevent this problem, enable rate limiting to ensure the ingress data rate does not exceed the egress rate of the ***slowest*** link in the IP WAN path. Rate limiting reduces the probability a device buffer will overflow, cause packet retransmits, and invoke TCP flow control. This sustains and maximizes WAN link throughput. Figure 4-14 illustrates rate limiting.

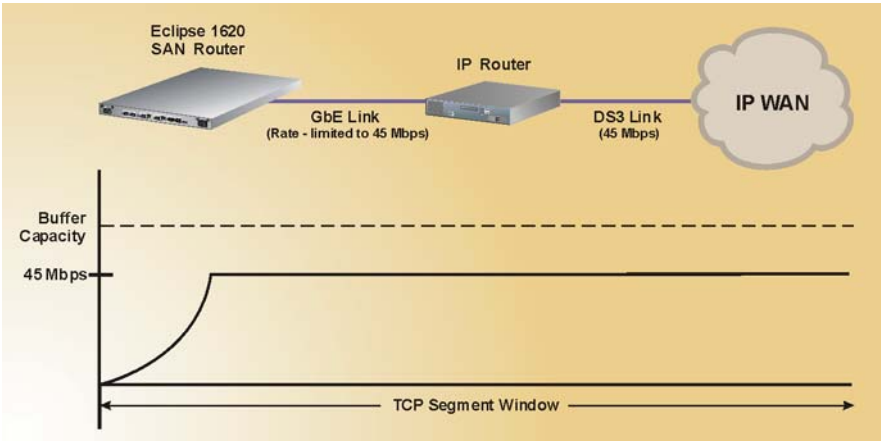

Figure 4-14 WAN Link Performance (Rate Limiting Enabled)

When configuring a SAN router for extended-distance operation over an IP WAN link, the peak available bandwidth must be determined or obtained from the network service provider, and storage traffic over the link must be rate-limited accordingly.

If the IP WAN link is dedicated, the peak available bandwidth equals the total link bandwidth. This implies that no other application data or traffic is routed across the link. If the IP WAN link is shared or channelized, the peak available bandwidth equals that portion of the total link bandwidth allotted for storage traffic at peak use time.

The speed of SAN router traffic (iFCP or iSCSI protocol) must be rate limited to the peak available bandwidth to prevent buffer overflow and dropped packets at intervening IP link networking equipment. Rate limiting is configured for SAN router ports as follows:

- Eclipse 1620 SAN Router - Two user-configured intelligent ports (**ETHERNET 3** and **4**) can be configured for IP network connectivity.
- Eclipse 2640 SAN Router - Four user-configurable intelligent ports (**13** through **16**) can be configured for IP network connectivity.

Rate limiting is set at the Element Manager application using the *Port Speed* drop-down list of the *FC/Ethernet Port Configuration* dialog box. The *Port Speed* drop-down list provides eight speed selections:

- **Digital Signal 1 (DS1) -** A framing and formatting specification that transmits 24 digital data channels on a T1 synchronous line. Each channel transmits at 64 Kbps (full-duplex), providing an aggregate bandwidth of 1.544 Mbps. Typical T1 lines are long-distance, point-to-point connections used for private networks and corporate Internet communication. The iFCP protocol overhead for a clear-channel (private) T1 link is 4.4%, which provides a bandwidth of **1.48 Mbps** (**0.18 MBps**) for storage traffic.
- **Thin Ethernet -** A transmission medium specified by IEEE 802.3 that carries information at 10 Mbps (full-duplex) in baseband form using low-cost (50-ohm type RG58) coaxial cable. The specification was developed to enable communication of LAN-connected computers. Thin Ethernet is also called 10 Base-T. The iFCP protocol overhead for a thin Ethernet link is 6.3%, which provides a bandwidth of **9.37 Mbps** (**1.17 MBps**) for storage traffic.
- **Digital Signal 3 (DS3) -** A framing and formatting specification that transmits 672 digital data channels on a T3 synchronous line. Each channel transmits at 64 Kbps (full-duplex), providing an aggregate bandwidth of 44.736 Mbps. Typical T3 lines are long-distance, point-to-point connections used by ISPs connecting to the Internet backbone and for the backbone itself.
  - The iFCP protocol overhead for a clear-channel (private) T3 link is 5.1%, which provides a bandwidth of **42.48 Mbps** (**5.31 MBps**) for storage traffic.
  - The iFCP protocol overhead for an asynchronous transfer mode (ATM) DS3 link is 22.7%, which provides a bandwidth of **34.59 Mbps** (**4.32 MBps**) for storage traffic.
- **Optical Container 1 (OC-1) -** A specification that defines the transport level for SONET traffic transmitted at 51.84 Mbps (full-duplex) using fiber-optic cable.
  - The iFCP protocol overhead for a packet over SONET (PoS) OC-1 link is 8.9%, which provides a bandwidth of **47.22 Mbps** (**5.90 MBps**) for storage traffic.
  - The iFCP protocol overhead for an ATM OC-1 link is 17.9%, which provides a bandwidth of **42.58 Mbps** (**5.32 MBps**) for storage traffic.

- **Fast Ethernet -** A transmission medium specified by IEEE 802.3 that carries information at 100 Mbps (full-duplex) in baseband form using Category-5 copper cable or fiber-optic cable. The specification was developed to enable faster communication of LAN-connected computers. Fast Ethernet is also called 100 Base-T. The iFCP protocol overhead for a fast Ethernet link is 6.3%, which provides a bandwidth of **93.70 Mbps** (**11.71 MBps**) for storage traffic.
- **Optical Container 3 (OC-3) -** A specification that defines the transport level for SONET traffic transmitted at 155.52 Mbps (full-duplex) using fiber-optic cable.
  - The iFCP protocol overhead for a PoS OC-3 link is 8.9%, which provides a bandwidth of **141.65 Mbps** (**17.71 MBps**) for storage traffic.
  - The iFCP protocol overhead for an ATM OC-3 link is 17.9%, which provides a bandwidth of **127.74 Mbps** (**15.97 MBps**) for storage traffic.
- **Optical Container 12 (OC-12) -** A specification that defines the transport level for SONET traffic transmitted at 622.08 Mbps (full-duplex) using fiber-optic cable.
  - The iFCP protocol overhead for a PoS OC-12 link is 8.9%, which provides a bandwidth of **566.59 Mbps** (**70.82 MBps**) for storage traffic.
  - The iFCP protocol overhead for an ATM OC-12 link is 17.9%, which provides a bandwidth of **510.98 Mbps** (**63.87 MBps**) for storage traffic.
- **Gigabit Ethernet -** A transmission medium specified by IEEE 802.3 that carries information at 1,000 Mbps (full-duplex) in baseband form using Category-5 copper or fiber-optic cable. The iFCP protocol overhead for a GbE link is 6.3%, which provides a bandwidth of **937.00 Mbps** (**117.12 MBps**) for storage traffic.

As a practical example, suppose the changes to a one-terabyte database must be backed up daily on a real-time basis. The changes constitute 10% of the database per eight-hour working day. Imposing a 2:1 data compression ratio and performing the computations yields a backup requirement of 1.74 MBps. An DS3 link (ATM) with a bandwidth of 4.32 MBps is the appropriate choice. This link provides nearly 2.5 times the required bandwidth to account for current storage traffic, unexpected burstiness, and capacity planning.

Distance Extension Best Practices

To implement a successful extended-distance BC / DR solution, follow a set of best practice conventions as follows:

1. **Use dedicated bandwidth and rate limiting -** If possible, negotiate dedicated bandwidth as part of the SLA with the network service provider. Enable intelligent port rate limiting to ensure the ingress data rate does not exceed the negotiated bandwidth.

   — **RDR/S applications -** Provision the bandwidth for peak traffic demand, unexpected bursts, and growth.

   — **RDR/A applications -** Provision the bandwidth for average traffic demand and growth.

   If dedicated bandwidth is not available, quality of service (QoS) processing applied to shared bandwidth may be acceptable. Ensure other applications using the shared bandwidth are characterized and understood. Best-effort shared bandwidth is not recommended.

2. **Optimize IP WAN use -** In addition to rate limiting, employ additional techniques to optimize the IP WAN. These include:

   — **Buffering -** To regulate data flow and smooth the inherent burstiness of storage traffic, enable a large (256 megabyte) transmit buffer for each long-link port.

   — **Flow control -** In conjunction with buffering, TCP and GbE provide flow control mechanisms.

   TCP provides sliding-window, end-to-end flow control at the transport layer (IP does not provide network layer flow control). However, the TCP flow control mechanism is inefficient and requires retransmits.

   If the IEEE 802.3x Ethernet flow control standard is enabled by GbE switches in an extended-distance link, SAN routers negotiate the use of flow control with these switches. Whenever possible, the best practice is to use IEEE 802.3x flow control to relieve input buffer congestion.

   — **Data compression -** Enable a compression algorithm to ensure data is compact and efficiently transmitted. The data compression ratio is a function of the data itself. Most data streams are compressible from between 2:1 and 15:1.

— **Jumbo frames -** To prevent fragmentation of Fibre Channel frames into multiple IP datagrams, enable jumbo frames to increase the data packet size from 1,500 bytes to approximately 9,000 bytes. Ensure the technology is supported by all IP network equipment in the data link.

— **FastWrite technology -** Enable FastWrite technology to reduce protocol overhead for extended-distance write transactions. The technology is very efficient over long distances with large write transactions (such as RDR/S applications).

— **Bandwidth management -** Enable QoS processing to guarantee bandwidth over a shared link. QoS subdivides port buffers into multiple queues, each with one or more associated drop thresholds. Multiple queues and drop thresholds allow the switch to prioritize output when faced with congestion.

3. **Minimize fabric hop count -** The maximum supported hop count in a fabric is three. Because the E_Port-to-R_Port ISL between a fabric element and SAN router counts as a hop, SAN routing connectivity is limited to two-hop fabrics. A remotely connected fabric does not add to the hop count of a local fabric; remote devices appear connected to the SAN router using proxy Domain_IDs 30 and 31. The best practice is to directly connect a SAN router to server and storage ports. However, SAN router connectivity through a fabric element is practical if the topology is required for scalability.

4. **Configure dual storage array controllers -** Storage device OEMs provide at least two controllers per storage array. Although RDR/S and RDR/A applications work with a single controller, use at least two controllers to provide high availability. Each controller has multiple Fibre Channel N_Ports that can be assigned to the RDR/S or RDR/A application. Most remote data replication software can load balance and initiate failover across the controllers.

5. **Implement parallel-path architecture -** It is recommended to configure redundant, parallel extended-distance links. It is also important to keep the links as homogeneous as possible. Some remote data replication applications are sensitive to path differences and decrease performance to the lowest common denominator. The best practice is to configure a dual-link architecture with similar paths (bandwidth and latency), through which RDR/S or RDR/A software performs load balancing and failover.

6. **Do not implement IP network failover -** Implement extended-distance link failover through the remote data replication software, not the IP network. Many RDR/S and RDR/A software OEMs do not support IP network link failover.
7. **Zone controller port pairs -** When implementing load sharing, create a separate zone for each pair of communicating ports (one initiator and one target per zone). Assign the zones to different intelligent (iFCP) ports on the SAN router. If the IP network correctly distributes the load across two paths, then load sharing is implemented. If the RDR/S or RDR/A application performs load balancing across two controllers, then load balancing is implemented. Perform zoning by node WWN. Port zoning is not supported between mSANs (through iFCP).
8. **Do not use default zoning -** Do not implement default zoning through the SANvergence Manager application. When performing an installation, the application displays a single zone named **zone1** (with corresponding **ID:1**). This zone is for default zoning and should not be used.
9. **Dedicate storage ports to remote data replication -** Storage array ports should be dedicated to RDR/S or RDR/A traffic. Do not share replication and local traffic applications on the same port.
10. **Set GbE switches to auto-negotiate -** SAN router GbE ports operate only at GbE speed, therefore GbE Ethernet switches should be set to auto-negotiate with SAN routers to provide connectivity. The setting can be disabled only if both devices are set to *not* auto-negotiate.
11. **Set data compression level to Auto -** Two problems associated with data compression are incorrect algorithm selection or compression which is ineffectual because the extended-link bandwidth exceeds the SAN router bandwidth. Set the compression level to **Auto** at the *Advanced TCP Configuration* dialog box (for any selected compression algorithm). Using this compression level mode, the SAN router automatically detects if compression can be used.
12. **Configure the iFCP timeout setting to 10 seconds -** At the *iFCP Setup* dialog box, set the default remote timeout setting (iFCP timeout) to 10 seconds for all RDR/A applications.

13. **Configure out-of-band product management -** SAN routers can be managed through inband or out-of-band connectivity. Inband management is provided through the same GbE connections used for iFCP storage traffic. Out-of-band management is provided through a data center LAN that connects servers, workstations, and other network-related equipment. Out-of-band product management is recommended.
14. **Explicitly assign unique mSAN_IDs -** When iSAN routing is implemented, each Eclipse 1620 SAN router comprises an mSAN (not the case with an Eclipse 2640 SAN router). Each mSAN must be assigned a unique mSAN_ID. The ID ranges from **0** to **255**, and is typically the last octet of the management port IP address (although not a requirement). Do not install a SAN router without changing the default mSAN_ID. Configure a unique ID for both SAN routers in an extended-distance link.
15. **Back up critical data -** Always back up SAN router access passwords, configurations, and zones. This avoids the possibility of having to reconfigure a SAN router as if it were a new installation.

Consolidating and Integrating iSCSI Servers and Storage

The iSCSI protocol defines rules and processes for transporting block-level small computer systems interface (SCSI) data over a TCP / IP network. iSCSI is designed as a protocol for an initiator to send SCSI commands to a target over IP.

iSCSI initiators (servers) include host bus adapters (HBAs) with iSCSI capability implemented in the hardware adapter card and software initiators running over standard network interface cards (NICs). iSCSI targets (storage) include disk storage systems, tape storage systems, and iSCSI gateways (such as Eclipse-series SAN routers).

Server HBAs and storage array NICs connect iSCSI resources over an IP network. Core transport layers are managed with existing network applications and high-level management activities of the iSCSI protocol (such as permissions, device information, and configuration) are layered over these applications.

The following sections describe:

- iSCSI protocol.
- iSCSI server consolidation.
- SCSI storage consolidation.

iSCSI Protocol

iSCSI is based on SCSI protocol that enables hosts to perform block data I/O operations with a variety of target peripherals. Targets include disk drives, tape devices, optical storage devices, printers, and scanners. A standard host-to-peripheral SCSI connection is based on a parallel transport mechanism with inherent distance and device support limitations. For storage applications, these limitations have caused development of high-speed serial transport technologies based on networking architectures such as Fibre Channel and GbE. IP storage networks based on serial gigabit transport layers overcome the distance, performance, scalability, and availability restrictions of parallel SCSI implementations.

By using SCSI protocols over network infrastructures, storage networking enables flexible, high-speed block data transfers for applications like tape backup, server clustering, storage consolidation, and disaster recovery.

iSCSI protocol defines a means to enable block storage applications over TCP/IP networks. An iSCSI initiator is typically a host (such as a file server) that issues requests to read or write data. The target is a passive resource (such as a disk array) that responds to initiator requests. When a server application sends a request, the operating system generates a packet with SCSI commands and a data request. The packet is encapsulated and encrypted (if required). A packet header is added and the resulting IP packet is transmitted over the TCP/IP network. The target storage device decrypts and disassembles the packet, then separates the SCSI commands and request. SCSI commands are transmitted to the SCSI controller, then to the SCSI storage device. Because iSCSI is bidirectional, the protocol returns data in response to the original request.

Compared to the standard SCSI protocol, Fibre Channel provides flexibility in terms of distance extension and switching capabilities. Fibre Channel also preserves the common SCSI controller application programming interface (API). Fibre Channel and iSCSI both preserve the SCSI command set. These common features allow deployment of storage solutions that rely on a combination of parallel SCSI and serial Fibre Channel technologies.

iSCSI Server Consolidation

Many enterprise-level IT departments have deployed decentralized computing configurations that include low-end, iSCSI-enabled servers directly attached to storage. While server acquisition costs are typically low, licensing and maintenance costs are often very high in terms of dollars and personnel time. The decentralized infrastructure also causes availability and reliability problems. For example:

- Many servers quickly run out of peripheral component interconnect (PCI) slots for adding direct-attached disk adapters. In addition, there is no room internal to the server to add hard drives. Therefore, more servers must be purchased.
- There is no ability to connect a server with ample disk space to another server with insufficient disk space. Efficient disk utilization is not possible.
- The failure rate of inexpensive external disks is high. In addition, server downtime must be planned to perform disk administration.
- Administrators work long hours (nights and weekends) to perform server maintenance, system updates, and other critical tasks. This leads to personnel dissatisfaction.

Server consolidation addresses cost, time, and availability issues by providing iSCSI-based server connectivity to a Fibre Channel storage fabric. Figure 4-15 illustrates iSCSI server consolidation.

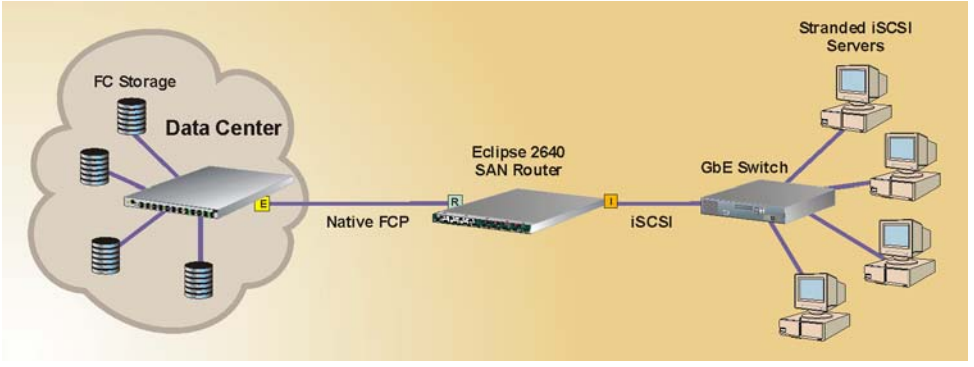

Figure 4-15 iSCSI Server Consolidation

As shown in the figure, server consolidation is enabled by installing an Eclipse 2640 SAN Router that provides iSCSI-to-native FCP connectivity.

Stranded servers subject to consolidation do not require installation at one physical location. Servers may be located in a data center or remotely; however, consolidation provides logical connectivity and access as though the servers were co-located. Servers can now access robust, scalable, and easily managed SAN storage that provides better data availability.

iSCSI Storage Consolidation

In general, the largest expense associated with an IT infrastructure storage is the purchase of storage (disk and tape drives). However, conventional architectures create multiple storage islands that make efficient disk utilization nearly impossible. Many environments have disk utilization rates below 50%.

Storage consolidation pools disk resources (no matter the location) and provides disk management as a single entity shared between servers. Consolidation addresses cost and utilization issues by providing iSCSI-based storage connectivity to Fibre Channel servers. Figure 4-16 illustrates iSCSI storage consolidation.

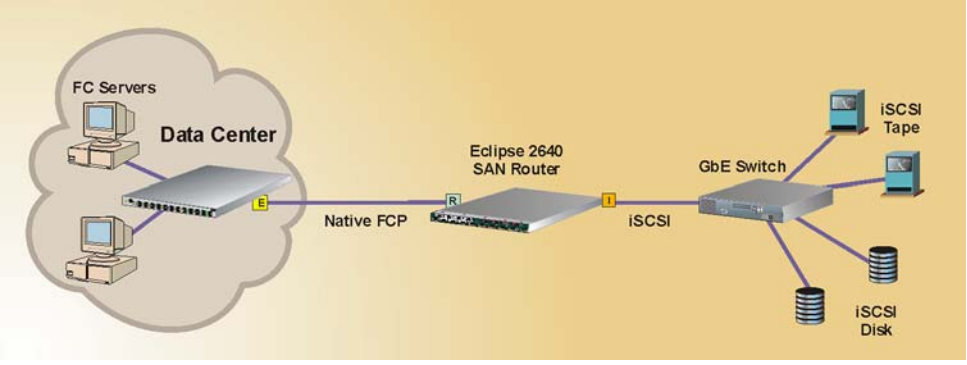

Figure 4-16 iSCSI Storage Consolidation

As shown in the figure, storage consolidation is enabled by installing an Eclipse 2640 SAN Router that provides iSCSI-to-native FCP connectivity.

5

Physical Planning Considerations

This chapter describes physical planning considerations for incorporating McDATA directors and switches into storage area networks (SANs) and Fibre Channel fabric topologies. The chapter provides planning considerations and recommendations for:

- Port connectivity and fiber-optic cabling.
- Rack-mount management server, Ethernet local area network (LAN), and remote access support.
- Security provisions for access to directors, switches, or the management server (password protection), and for customer data paths through directors, fabric switches, and SAN routers.
- Optional feature keys.

Port Connectivity and Fiber-Optic Cabling

This section provides planning recommendations for:

- Port requirements (number, type, and speed of ports).
- Small form factor pluggable (SFP) optical transceivers.
- Extended-distance ports.
- High-availability considerations.
- Cabling and connectors.
- Routing fiber-optic cables.

Port Requirements

Plan for sufficient shortwave laser, longwave laser, and 1.0625, 2.1250, 4.2500, and 10.2000 gigabit per second (Gbps) Fibre Channel ports to meet the needs of the SAN configuration. The number of ports required is equal to the number of device connections (including redundant connections), plus the number of interswitch links (ISLs) between fabric elements, plus the total number of spare port connections.

The number of Fibre Channel ports and port operation for McDATA directors and switches are described as follows:

- **Intrepid 6064 Director -** The director is configured from a minimum of eight fiber port module (FPM), universal port module (UPM), or ten-gigabit port module (XPM) cards (32 ports total) to a maximum of 16 cards (64 ports total).
  - — FPM cards provide four 1.0625 Gbps port connections and can be configured with a combination of shortwave or longwave transceivers.
  - — UPM cards provide four 2.1250 Gbps port connections and can be configured with a combination of shortwave or longwave transceivers.
  - — XPM cards provide one 10.2000 Gbps port connection and can be configured with shortwave or longwave transceivers.
- **Intrepid 6140 Director -** The director is configured from a minimum of 16 UPM or XPM cards (64 ports total) to a maximum of 35 cards (140 ports total).
  - — UPM cards provide four 2.1250 Gbps port connections and can be configured with a combination of shortwave or longwave transceivers.
  - — XPM cards provide one 10.2000 Gbps port connection and can be configured with shortwave or longwave transceivers.
- **Intrepid 10000 Director -** The director is configured from a minimum of one to a maximum of eight line modules (LIMs). Each LIM provides the interface to attach up to four optical paddles as follows:
  - — Optical paddles that operate at 1.0625 or 2.1250 Gbps provide eight Fibre Channel port connections. A fully-populated director supports up to 256 connections and can be configured with a combination of shortwave or longwave transceivers.

— Optical paddles that operate at 10.2000 Gbps provide two Fibre Channel port connections. A fully-populated director supports up to 64 connections and can be configured with a combination of shortwave or longwave transceivers.

- **Sphereon 3232 Fabric Switch -** The switch provides up to 32 duplex SFP fiber-optic port transceivers (1.0625 or 2.1250 Gbps operation). Shortwave laser and longwave laser transceivers are available.
- **Sphereon 4300 Fabric Switch -** The switch provides up to 12 duplex SFP fiber-optic port transceivers (1.0625 or 2.1250 Gbps operation). Shortwave laser and longwave laser transceivers are available.
- **Sphereon 4400 Fabric Switch -** The switch provides up to 16 duplex SFP fiber-optic port transceivers (1.0625, 2.1250, or 4.2500 Gbps operation). Shortwave laser and longwave laser transceivers are available.
- **Sphereon 4500 Fabric Switch -** The switch provides up to 24 duplex SFP fiber-optic port transceivers (1.0625 or 2.1250 Gbps operation). Shortwave laser and longwave laser transceivers are available.
- **Sphereon 4700 Fabric Switch -** The switch provides up to 32 duplex SFP fiber-optic port transceivers (1.0625, 2.1250, or 4.2500 Gbps operation). Shortwave laser and longwave laser transceivers are available.
- **Eclipse 1620 SAN Router -**Two user-configured Fibre Channel ports provide 1.0625 Gbps connectivity using SFP port connectors and two user-configured intelligent ports provide Fibre Channel and Internet protocol (IP) network connectivity. Each intelligent port provides two connectors (SFP or RJ-45).
- **Eclipse 2640 SAN Router -** Twelve user-configured ports provide 1.0625 or 2.1250 Gbps Fibre Channel or user datagram protocol (UDP) connectivity using SFP port connectors. Four user-configurable intelligent ports provide Internet Fibre Channel protocol (iFCP) or Internet small computer systems interface (iSCSI) network connectivity using SFP port connectors.

SFP Optical Transceivers

Shortwave laser SFP optical transceivers (1.0625, 2.1250, 4.2500, or 10.2000 Gbps) provide a connection for multimode cable with a core diameter of 50 microns and a cladding diameter of 125 microns (50/125 micron), or multimode cable with a core diameter of 62.50 microns and a cladding diameter of 125 microns (62.5/125 micron).

Longwave laser SFP optical transceivers (1.0625, 2.1250, 4.2500, or 10.2000 Gbps) provide a connection for singlemode cable with a core diameter of 9 microns and a cladding diameter of 125 microns (9/125 micron).

Consider the following when determining the number and type of transceivers to use:

- Distance between a director or fabric switch and the attached Fibre Channel device or between fabric elements communicating through an ISL.
- Cost effectiveness.
- Device restrictions or requirements with respect to existing fiber-optic (multimode or singlemode) or copper cable.

Data Transmission Distance

Data transmission distance is a factor governing the choice of transceiver type, fiber-optic cable type, and transmission rate. When using multimode cable, if the core diameter or data transmission rate increases, the data transmission distance decreases.

Link budget is another governing factor. A link budget is the attenuation (in dB) a connection between devices can sustain before significant link errors or loss of signal occur. When using multimode cable, if the core diameter or data transmission rate increases, the link budget decreases.

Cable-conversion, repeater, patch-panel, or other connections within a link also decrease the link budget. Each connection introduces a nominal signal loss of at least one dB through the link. Patch panel connections (with one connection at each side of the panel) typically introduce a two dB signal loss through a link.

Other variables such as the grade of fiber-optic cable, device restrictions, application restrictions, buffer-to-buffer credit limits, and performance requirements can also affect data transmission distance and link budget.

Table 5-1 lists unrepeated data transmission distance and link budget as a function of fiber-optic cable type and data transmission rate.

When using multimode cable, note the decrease in performance as the cable core diameter or data transmission rate increases. When using singlemode cable, performance is a function of transceiver type. Data transmission distance and link budget are not affected by data transmission rate.

Table 5-1 Cable Type and Transmission Rate versus Distance and Link Budget

Cable Type and Data Transmission Rate	Unrepeated Distance	Link Budget
62.5/125 micron multimode at 1.0625 Gbps	300 meters (984 feet)	3.0 dB
62.5/125 micron multimode at 2.1250 Gbps	150 meters (492 feet)	2.2 dB
62.5/125 micron multimode at 4.2500 Gbps	75 meters (246 feet)	2.3 dB
62.5/125 micron multimode at 10.2000 Gbps	33 meters (108 feet)	2.4 dB
50/125 micron multimode at 1.0625 Gbps	500 meters (1,640 feet)	3.9 dB
50/125 micron multimode at 2.1250 Gbps	300 meters (984 feet)	2.8 dB
50/125 micron multimode at 4.2500 Gbps	150 meters (492 feet)	2.5 dB
50/125 micron multimode at 10.2000 Gbps	82 meters (269 feet)	2.2 dB
9/125 micron singlemode at 1.0625, 2.1250, 4.2500, or 10.2000 Gbps (10-kilometer SFP optical transceiver)	10.0 kilometers (6.2 miles)	7.8 dB
9/125 micron singlemode at 1.0625, 2.1250, 4.2500, or 10.2000 Gbps (20-kilometer SFP optical transceiver)	20.0 kilometers (12.4 miles)	7.8 dB
9/125 micron singlemode at 1.0625, 2.1250, 4.2500, or 10.2000 Gbps (35-kilometer SFP optical transceiver)	35.0 kilometers (21.7 miles)	7.8 dB

Cost Effectiveness

Cost is another factor governing the choice of transceiver type and optical fiber. Shortwave laser transceivers and multimode cable offer a less expensive solution if data transmission distance is not critical.

Device or Cable Restrictions

The choice of transceiver and cable type may be restricted or dictated by:

- **Device restrictions -** Some devices may be restricted to use of only one type of transceiver (shortwave or longwave). Refer to the device's supporting documentation for information.

- **Existing cable restrictions -** The enterprise may contain only one type of fiber-optic cable (multimode or singlemode), and the customer may be required to use the existing cables. Customers may also be required to use existing copper cables for some arbitrated loop devices.

Extended-Distance Ports

Through longwave laser transceivers and repeaters or wavelength division multiplexing (WDM) equipment, directors and fabric switches support Fibre Channel data transmission distances of over 100 km. The extended distance feature is enabled on a port-by-port basis by using entries in the *RX BB Credit* column for a specified port at the Element Manager application's *Configure Ports* dialog box. This feature provides extended distance support using Fibre Channel protocol. Refer to *Distance Extension Through BB_Credit* for additional information.

When a director or fabric switch port is configured to support extended link distances, the attached device (or attached fabric element) must also support extended distance operation and be configured to use a higher BB_Credit value to maintain link efficiency. If the extended distance feature is enabled for a port that is not installed or does not support extended distance operation, the configuration for the feature is ignored.

High-Availability Considerations

To provide high device availability, critical servers, storage devices, or applications should be connected to more than one fabric element (director or switch) or to more than one fabric. To determine if dual-connection capability exists for a device, refer to the associated device documentation. To provide high fabric availability, consider the use of multiple fabric elements (directors and switches), multiple ISLs, or redundant fabrics. Refer to *Fabric Availability* for information.

Plan to maintain unused (spare) director and switch ports if port connections must be quickly moved and re-established after a failure. If an individual port or an entire port card fails, optical transceivers or port cards can be removed and replaced, spare port connections identified (through the Element Manager application), and fiber-optic cables rerouted and reconnected while the director or switch is operational.

Fibre Channel Cables and Connectors

This section provides Fibre Channel cable and connector planning information as follows:

- Cables for directors, fabric switches, and SAN routers.
- Intrepid-series director, Sphereon-series fabric switch, and Eclipse-series SAN router optical connectors.

Cables

Fiber-optic jumper cables are required to connect directors, fabric switches, and SAN routers ports to servers, devices, distribution panels, or other elements in a multiswitch fabric or routed SAN. Depending on the attached device and fabric element port, use one of the following types of cable:

- Graded-index 62.5/125 micron multimode cable provides a transmission distance of up to 300 meters (1.0625 Gbps), 150 meters (2.1250 Gbps), 75 meters (4.2500 Gbps), or 33 meters (10.2000 Gbps) and connects to shortwave ports that transmit light at an 850 nanometer (nm) wavelength. The cable typically has an orange jacket.
- Graded-index 50/125 micron multimode cable provides a transmission distance of up to 500 meters (1.0625 Gbps), 300 meters (2.1250 Gbps), 150 meters (4.2500 Gbps), or 82 meters (10.2000 Gbps) and connects to shortwave ports that transmit light at an 850 nm wavelength. The cable typically has an orange jacket.
- Depending on transceiver type, dispersion-unshifted (step-index) 9/125 micron singlemode cable provides a transmission distance of up to 10, 20, or 35 kilometers and connects to longwave ports that transmit light at a 1300 nm wavelength. The cable typically has a yellow jacket.

LC Connectors

Multimode or singlemode cables attach to Intrepid-series director, Sphereon-series fabric switch, and Eclipse-series SAN router ports with SFP optical transceivers and LC duplex connectors. Figure 5-1 illustrates an SFP transceiver and LC duplex connector.

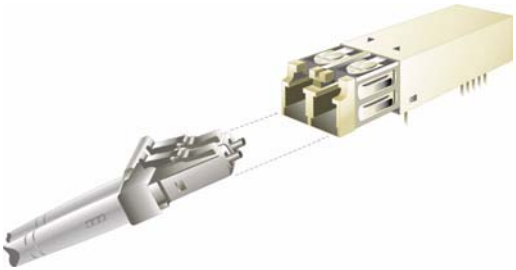

Figure 5-1 SFP Transceiver and LC Duplex Connector

Routing Fiber-Optic Cables

Follow a logical plan for routing fiber-optic cables to avoid confusing connections during installation and operation. Route cables from the access holes at the bottom of the Fabricenter equipment cabinet to fabric element ports. When routing cables to ports be aware:

- In a Fibre Channel Protocol (FCP) environment, ports are numbered by physical port number.
- In a fibre connection (FICON) environment, ports are numbered by logical port address. The translation between physical port number and logical port address varies by equipment type and original equipment manufacturer (OEM).

Figure 3-17 and Figure 3-18 illustrate port numbering and logical port addressing for the Intrepid 6140 Director. Although the figures depict a UPM card map only for the Intrepid 6140 Director, physical port numbers and logical port addresses can be extrapolated for other switch products.

Leave enough slack in the cables to allow cable movement for FPM card, UPM card, XPM card, LIM, or SFP optical transceiver removal and replacement or possible rerouting of the cable to another port. After cables are routed and connected, secure the cables to the sides of the cabinet using cable ties provided. When routing fiber-optic cables and estimating cable lengths, consider:

- Cable routing inside the equipment cabinet to different port locations, and installation position of the director or switch (top or bottom of the cabinet). Plan for 1.0 meter (39.37 inches) of extra cable for routing through restraint mechanisms and rerouting cables to other ports.

- Cable routing outside the equipment cabinet. Plan for 1.5 meters (5 feet) of cable outside the cabinet to provide slack for service clearance, limited cabinet movement, and inadvertent cable pulls.
- Cabling distance to servers, storage devices, and directors (for multiswitch fabric support).

The need for additional fiber-optic cabling could grow rapidly. More cables may be required for connections to additional servers or storage devices, or for connections to additional fabric elements as a multiswitch fabric is developed. The director or switch may need to be moved for more efficient connection to other units but still maintain its original connections. To account for these possibilities, consider installing excess fiber-optic cables.

Management Server, LAN, and Remote Access Support

Out-of-band (non-Fibre Channel) console access to directors, fabric switches, and SAN routers is provided to perform a variety of operations and management functions. These functions are performed from one or more of the following consoles:

- Through a personal computer (PC) or workstation connected to the management server through a customer LAN segment. The server is LAN-attached to the Ethernet port on a director control processor (CTP) card, fabric switch front panel, or SAN router front panel.
- Through a simple network management protocol (SNMP) management workstation connected through the director, fabric switch, or SAN router LAN segment; or the customer intranet.
- Through a PC with a web browser and Internet connection to the Enterprise Fabric Connectivity Manager (EFCM) Basic Edition interface on the director or fabric switch.
- Through a PC with a direct serial connection to the director, fabric switch, or SAN router maintenance port. The maintenance port is used by installation personnel to configure product network addresses.
- Through a PC with a modem connection to the management server. The modem is for use by support center personnel only.

Management Server

The management server is rack-mounted in a Fabricenter equipment cabinet. The server supports up to 48 McDATA directors, fabric switches, or SAN routers (managed products). The server is used to configure the product and associated SAN management and Element Manager applications, monitor product operation, change configurations, download firmware updates, and initiate diagnostics.

NOTE: The Sphereon 4300 Fabric Switch is not supported by the management server.

A server failure does not affect port connections or functions of an operational fabric element. The only operating effect of a server failure is loss of remote access, configuration, management, and monitoring functions.

Management Server Connectivity

The management server provides an auto-detecting 10/100 Base-T Ethernet interface that connects to the 24-port hub mounted at the top of the Fabricenter equipment cabinet. Each director CTP card, fabric switch front panel, or SAN router front panel also provides an auto-detecting 10/100 Base-T Ethernet interface that connects to the hub. Factory-installed cables connect the management server, hub, and managed products.

Although directors provide two Ethernet connections to the hub, only one connection is active at a time. The interface on the backup CTP card remains passive until a failure on the active CTP card occurs, at which point the redundant CTP card becomes active using the same media access control (MAC) address as the original interface.

If an optional customer intranet is used for LAN connections, the management server provides a second auto-detecting 10/100 Base-T Ethernet connection. This interface is used for remote workstation access.

The management server has an internal modem for service and support of managed products. The modem provides a dial-in capability that allows authorized service personnel to communicate with the management server and operate the SAN management and Element Manager applications remotely. The modem is also used to automatically dial out to an authorized support center (to report the occurrence of significant system events) using a call-home feature. The call-home feature is enabled in the Element Manager application and configured through the dial-up networking feature of Windows.

Connectivity Planning Considerations

Directors, fabric switches, SAN routers, and the management server are delivered in a cabinet-mount configuration in accordance with customer specifications. Because Ethernet cables that connect the managed products, hub, and management server are factory-installed, connectivity planning is not required for a stand-alone cabinet installation. However, consider the following Ethernet connectivity issues when:

- **Installing additional cabinet-mount products -** When installing an additional fabric element, the length of Ethernet cable required to provide hub connectivity is a function of cabinet position (top, bottom, or adjacent to the management server). Ensure cable lengths provide sufficient cable inside the cabinet to route to product Ethernet ports and to allow service clearance.
- **Interconnecting Fabricenter cabinets -** To increase the products managed by one management server, Ethernet hubs in one or more equipment cabinets must be connected. Plan for an Ethernet cable length that meets the distance requirement between cabinets. In addition, plan for an additional 1.5 meters (5 feet) of cable outside the cabinet to provide slack for service clearance, limited cabinet movement, or inadvertent cable pulls. Store extra Ethernet cable in the cabinet or under the computer room raised floor.
- **Consolidating management server operation -** For control and efficiency, all directors, fabric switches, and SAN routers in a multiswitch fabric or routed SAN should be managed by one management server. When products in two or more cabinets are joined to form a fabric, the PC environment should be consolidated to one server and one or more clients. Plan for Ethernet cabling to interconnect cabinets and ensure all fabric elements and PC platforms participating in the fabric have unique IP addresses.

Remote User Workstations

Customer system administrators determine whether to allow access to directors and switches from remote workstations. If administrators allow remote sessions, they may restrict access to selected workstations by configuring the IP addresses of those workstations through the SAN management application. When a remote session is allowed, the remote user has the same rights and permissions as if the session were on the local management server. Up to 25 sessions can be simultaneously active.

NOTE: Remote workstation access to Eclipse-series SAN routers is not supported.

Remote workstations must have access to the LAN segment on which the management server is installed. Product administrative functions are accessed through the LAN and management server. The LAN interface can be:

- Part of the dedicated 10/100 megabit per second (Mbps) LAN segment that provides access to managed products. This Ethernet connection is part of the equipment cabinet installation and is required. Connection of remote workstations through the hub is optional. This type of network configuration using one Ethernet connection through the management server is shown in Figure 5-2. Intrepid 6064 Directors are used as an example.

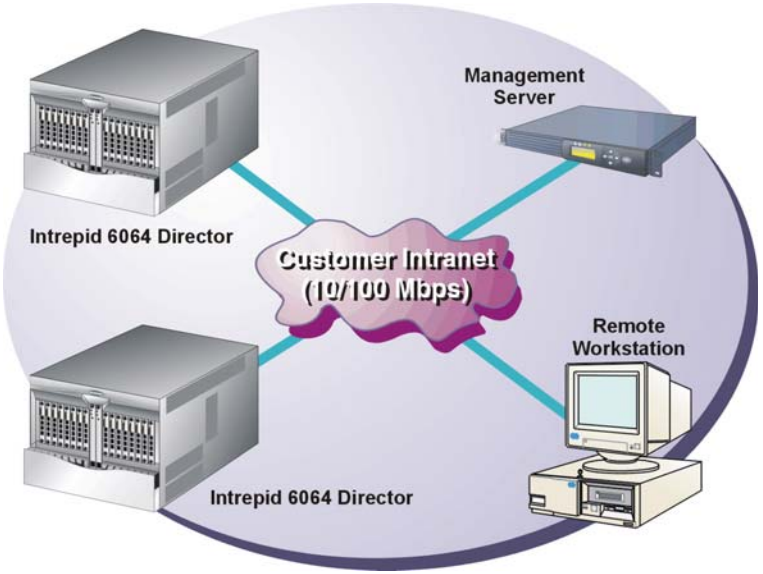

Figure 5-2 Typical Network Configuration (One Ethernet Connection)

- Part of a second management server interface that connects to a customer intranet and allows operation of the Element Manager application from remote user PCs or workstations. The customer intranet can be a ten or 100 Mbps LAN segment.

Connection to this LAN segment is optional and depends on customer requirements. This type of network configuration using both Ethernet connections is shown in Figure 5-3. Intrepid 6064 Directors are used as an example.

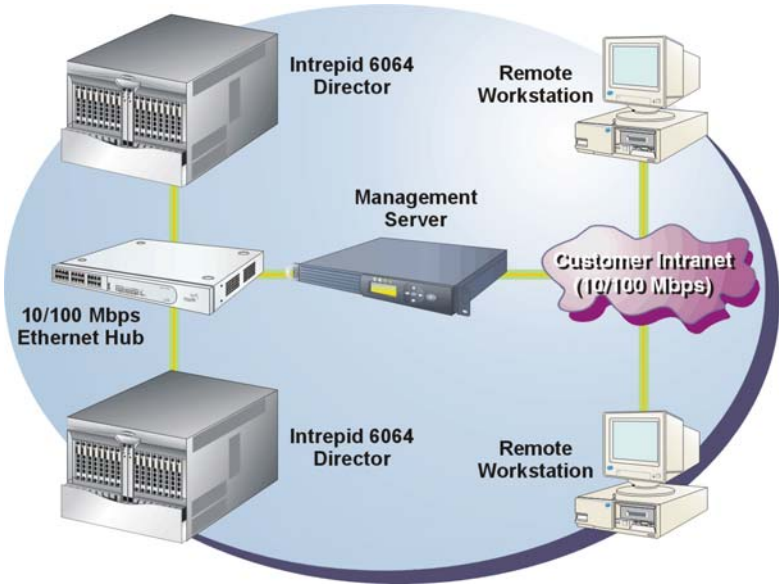

Figure 5-3 Typical Network Configuration (Two Ethernet Connections)

If only one management server connection is used and this connection is provided through the customer intranet, all functions provided by the server are available to users throughout the enterprise. The purpose for dual LAN connections is to provide a dedicated LAN segment that isolates the server and managed products from other users in the enterprise.

NOTE: Both Ethernet adapters in the management server provide auto-detecting 10/100 Mbps connections. The dedicated LAN segment that connects the server to managed products and the optional customer intranet operate at either ten or 100 Mbps.

SNMP Management Workstations

An SNMP agent that runs on the management server can be configured through the SAN management application. This agent implements Version 3.1 of the Fibre Alliance management information base (MIB) as follows:

- Up to 12 authorized management workstations can be configured through the director or fabric switch SAN management application to receive unsolicited SNMP trap messages that indicate product operational state changes and failure conditions.
- Up to eight authorized management workstations can be configured through a SAN router Element Manager application to send SNMP trap messages that indicate product operational state changes and failure conditions.

In addition, there is a separate SNMP agent that runs on each director, fabric switch, or SAN router (configured through the associated Element Manager application). The director or switch SNMP agent can be configured to send unsolicited SNMP trap messages to up to six recipients. The SAN router SNMP agent can be configured to send unsolicited SNMP trap messages to up to four recipients.

SNMP management is only intended for product monitoring; therefore, the default state of all MIB variables is read-only. If installed on a dedicated LAN, SNMP management workstations communicate directly with all managed products. If installed on a customer intranet, workstations communicate with managed products through the management server.

EFCM Basic Edition Interface

The EFCM Basic Edition interface provides a graphical user interface (GUI) accessed through the Internet (locally or remotely) to manage a single director or switch. If the EFCM Basic Edition interface is to be implemented:

- Plan for an Internet connection to the LAN segment on which the product is installed. The LAN connection is provided through the optional McDATA-supplied Ethernet hub or the corporate intranet.
- Ensure adequate security measures are implemented to preclude unauthorized access to managed products. Ensure IP addresses (uniform resource locators (URLs) for Internet access) of managed products, usernames, and passwords are tightly controlled.

NOTE: EFCM Basic Edition interface access to the Intrepid 10000 Director and SAN routers is not supported.

Security Provisions

Security provisions are available to restrict unauthorized access to a director, switch, or attached Fibre Channel devices. Access to the director or switch (through the SAN management application, Element Manager application, or EFCM Basic Edition interface) is restricted by implementing password protection. Access to attached computing resources (including applications and data) is restricted by implementing one or more of the following security provisions:

- SANtegrity Authentication.
- SANtegrity Binding.
- Prohibit dynamic connectivity mask (PDCM) arrays.
- Preferred path.
- Zoning.
- Server and storage-level access control.

Password Protection

Access to the SAN management and Element Manager applications requires configuration of a user name and password. Up to 16 user names and associated passwords can be configured. Each user is assigned rights that allow access to specific sets of product management operations. Table 5-2 explains the types of user rights available. A user may have more than one set of user rights granted.

Table 5-2 Types of User Rights

User Right	Operator Access Allowed
View Only	The user may view product configurations and status but may not make changes. These rights are the default if no other user rights are assigned.
Operator	The operator may view status and configuration information through the Element Manager application and perform operational control changes such as blocking ports and placing the product online or offline.
Product Administrator	The product administrator can make control and configuration changes through the Element Manager application.
System Administrator	The system administrator can make control and configuration changes, define users and passwords, and add or remove products through the SAN management application.
Maintenance	The maintenance operator can perform product control and configuration changes through the Element Manager application and perform diagnostics, maintenance functions, firmware loads, and data collection.

System administrators can use the SAN management application to assign remote workstation access to directors and switches. Remote sessions are allowed for anyone on a customer intranet, disallowed completely, or restricted to specific workstations. Remote users must log into the SAN management application with a user name and password, just as when logging in to the local management server. Passwords are encrypted when sent across the network. By entering workstation IP addresses at the SAN management application, administrators can allow access from all user workstations or from only specific workstations.

For access through the EFCM Basic Edition interface, the system administrator provides IP addresses of products to authorized users, assigns access usernames, and controls associated passwords.

SANtegrity Authentication

SANtegrity Authentication enhances SAN security by providing a set of user-configurable, software-enforced features that restrict access to Fibre Channel fabric elements. Features protect against accidental or intentional attacks to fabric elements by not allowing connection of devices or management interfaces that cannot be identified. Security features are independent from one another and may be individually enabled or disabled by an administrator. SANtegrity Authentication features include:

- **Password safety -** When accessing a director or fabric switch for the first time through the command line interface (CLI) or EFCM Basic Edition interface, the password must be changed. When accessing a director or switch for the first time through the maintenance port (enhanced serial authentication enabled), the password must be changed.

  Upon user login, the password is checked against the original default password. If the password and default password match, the user must change the password. This functionality addresses a common security defect where the default password is never changed.

- **Management server CHAP authentication -** Enhanced login security between a fabric element (director, fabric switch, or SAN router) and the management server is provided through challenge handshake authentication protocol (CHAP). A fabric element uses CHAP to authenticate any management server that attempts a connection.

The fabric element transmits a random value (used only once), an ID value (incremented at each login), and a shared CHAP secret (16-byte random value) to the server. The server concatenates the random value, ID value, and CHAP secret, and calculates a one-way message digest (also called a hash value). The hash value is transmitted to the authenticator (fabric element). The fabric element then builds the same concatenated string and compares the result with the value received from the server. If the values match, the connection is authenticated.

- **Port DHCHAP authentication -** Enhanced security for device connections and ISLs is provided through Diffie-Hellman challenge handshake authentication protocol (DHCHAP). A fabric element uses DHCHAP to authenticate any device (node) that attempts a node port (N_Port) connection and any director or switch that attempts an expansion port (E_Port) connection. This ensures only authorized devices can be added to the fabric.

  DHCHAP is an authentication protocol based on transmission of a one-way hash value (comprised of a sequentially-incremented ID value and CHAP secret). Because the hash cannot be reversed to discover the CHAP secret, the protocol provides protection from discovery through the network.

- **CT authentication -** Common transport (CT) authentication authorizes management server access to fabric elements through the open-system management server (OSMS) interface. The feature is software-enforced and allows an attached fabric to authenticate the OSMS management application. A single shared secret is configured for each fabric-attached director or switch (because OSMS is a fabric service that assumes all attached fabric elements are authenticated). The same secret is used by the management application.

- **PCP user database -** All authentication users are configured in a product control point (PCP) user database. The database includes usernames, passwords, and authorized interfaces for management server and device access. The database controls password authentication for EFCM, SANavigator, CLI, and EFCM Basic Edition management interfaces. The database also controls CHAP and CT authentication for Fibre Channel ports.

- **RADIUS server support -** Remote authentication dial-in user service (RADIUS) is a client-server, UDP-based protocol that supports storage and authentication of passwords and CHAP secrets. Directors, fabric switches, and SAN routers support a RADIUS client (LAN-connected to a primary or secondary RADIUS server) that authenticates CHAP responses and login passwords. The RADIUS server stores:
  - — Management server-to-fabric element (director or fabric switch) CHAP secrets.
  - — E_Port and N_Port DHCHAP secrets.
  - — Hypertext transfer protocol (HTTP) user passwords for the EFCM Basic Edition interface.
  - — Telnet user passwords for the CLI.
  - — RADIUS server interface encryption keys.
- **Inband access control list -** The management server interface supports an access control list (ACL) that provides attached port worldwide names (WWNs) or switch node names for which director or fabric switch communication is allowed. The CLI and EFCM Basic Edition interface do not support configuration of an inband access control list.
- **Out-of-band access control list -** Directors and fabric switches support an IP-based ACL that defines the node IP addresses that are permitted to log in to the fabric element through an out-of-band management interface. Each director or fabric switch is individually configured with a list of IP address ranges.
- **Encrypted SSH protocol -** Secure shell (SSH) protocol is a software-enforced security encryption feature that controls CLI access to a director or fabric switch. The SSH protocol suite supports secure shell communication, remote file copy, file transfer, and port forwarding through a telnet interface.
- **Security log -** The security log records security-related events (including but not limited to SANtegrity features). The log is a default feature of the Enterprise Operating System, classic (E/OSc) firmware and does not require enablement through a product feature enablement (PFE) key. Log entries record the following events:

— Authorization errors.

— Authentication errors.

— Management application user connections.

Use of the SANtegrity Authentication feature in conjunction with other security provisions must be carefully planned and coordinated. For additional information, refer to *Security Best Practices*. Obtain planning assistance from McDATA's professional services organization before implementing the feature.

SANtegrity Binding

SANtegrity Binding is a feature that enhances data security in large and complex SANs comprised of numerous fabrics and devices provided by multiple OEMs, SANs that intermix FCP and FICON protocols, and FICON-cascaded high-integrity SANs. The feature allows or prohibits director or switch attachment to fabrics (fabric binding) and Fibre Channel device attachment to directors or switches (switch binding). The SANtegrity Binding feature includes:

- **Fabric binding -** Using fabric binding, administrators allow only specified directors or fabric switches to attach to specified fabrics in a SAN. This provides security from accidental fabric merges or disruption, particularly in environments that use patch panels for centralizing fibers and physical connections. This feature is enabled through the SAN management application.
- **Switch binding -** Using switch binding, administrators allow only specified devices and fabric elements to connect to specified director or fabric switch ports. This provides security in environments that include a large number of devices by ensuring only the intended set of devices attach to a director or switch. This feature is enabled through the Element Manager application.

Enterprise Fabric Mode

Although *Enterprise Fabric Mode* is not a keyed feature, it is integral to SANtegrity Binding operation. *Enterprise Fabric Mode* must be enabled through the SAN management application before fabric binding and switch binding can operate. *Enterprise Fabric Mode* also enables the following parameters:

- **Rerouting delay -** If a fabric topology changes, directors and fabric switches calculate a new least-cost data transfer path through a fabric, and routing tables immediately implement that path. This may result in Fibre Channel frames being delivered to a destination device out of order, because frames transmitted over the new (shorter) path may arrive ahead of previously-transmitted frames that traverse the old (longer) path. When enabled, the rerouting delay parameter ensures frames are delivered through a fabric in the correct order.
- **Domain RSCNs -** Domain registered state change notifications (RSCNs) provide connectivity information to all host bus adapters (HBAs) and storage devices attached to a fabric. RSCNs are transmitted to all registered device N_Ports attached to a fabric if either a fabric-wide event or zoning configuration change occurs.
- **Insistent Domain_ID -** When this parameter is enabled, the domain identification (Domain_ID) configured as the preferred Domain_ID for a fabric element becomes the active Domain_ID when the fabric initializes. A static and unique active domain identification is required by the fabric binding feature because the feature's fabric membership list identifies fabric elements by WWN and Domain_ID. If a duplicate preferred Domain_ID is used, then insisted upon, a warning occurs and the affected director or fabric switch cannot be added to the membership list.

SANtegrity Binding Planning Considerations

Fabric and switch binding enhance data security by controlling and monitoring director, fabric switch, and device connectivity. In fact, installation of the SANtegrity Binding feature is a prerequisite for configuring a high-integrity, FICON-cascaded SAN.

Use of the SANtegrity Binding feature in conjunction with other security provisions must be carefully planned and coordinated. For additional information, refer to *Security Best Practices*. Obtain planning assistance from McDATA's professional services organization before implementing the feature.

PDCM Arrays

PDCM connectivity control is configured and managed at the director or fabric switch level using the *Configure Allow/Prohibit Matrix - Active* dialog box (Figure 5-4), where the user specifies an array in which logical port addresses are allowed or prohibited from connecting with each other (including E_Port connectivity).

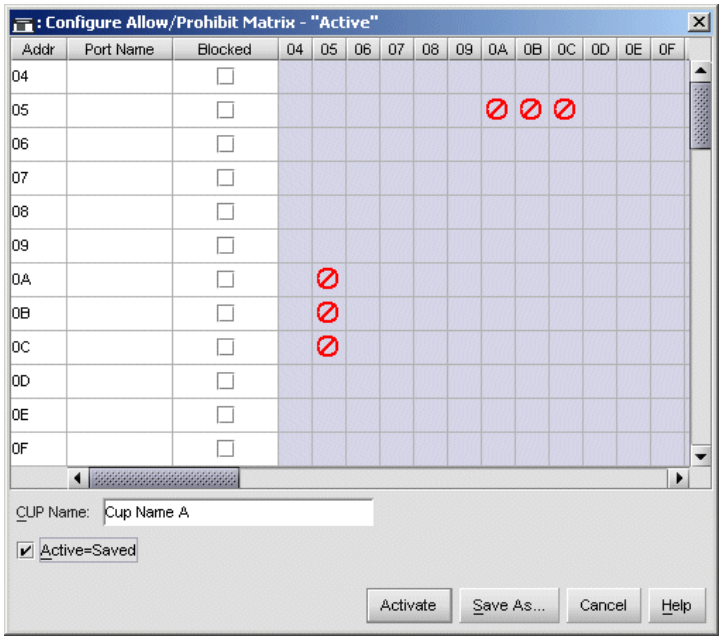

Figure 5-4 Configure Allow/Prohibit Matrix - Active Dialog Box

To access the dialog box, ensure the FICON management style is enabled for the director or switch, then select the *Allow/Prohibit* and *Active* options from the Element Manager application's *Configure* menu.

Figure 5-4 shows that port **1** (logical port address **05**) is prohibited from communicating with port **6** (logical port address **0A**), port **7** (logical port address **0B**), and port **8** (logical port address **0C**).

When implementing an array that prohibits E_Port connectivity, be aware that ISLs can be configured as unavailable to attached devices, causing complex routing problems that can be difficult to fault isolate and be incorrectly diagnosed as issues associated with the devices.

As an example of such a problem, refer to the simple two-director fabric illustrated in Figure 5-5. As shown in the figure, ISL **1** connects Director **A** and Director **B** through logical port addresses **09** and **1A**. ISL **2** connects the directors through logical port addresses **0A** and **1B**. A source server attaches to Director **A** through logical port address **05**. Two destination devices attach to Director **B** through logical port addresses **2C** and **2D**.

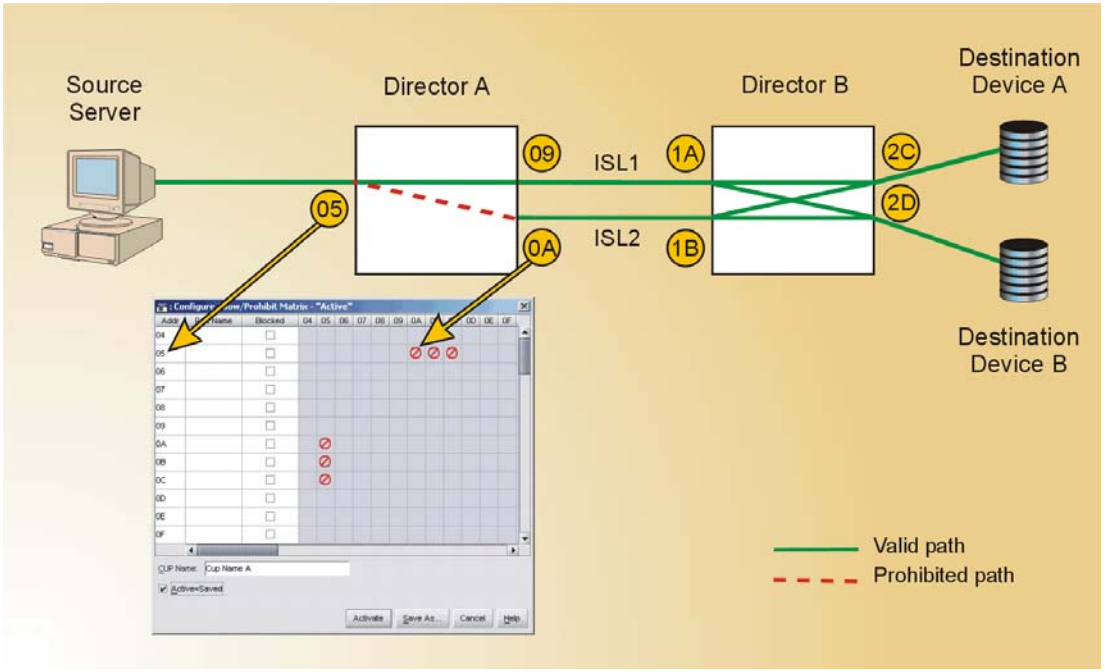

Figure 5-5 PDCM Array - Example Problem

A PDCM array configured for Director **A** prohibits logical port address **05** from communicating with logical port addresses **0A**, **0B**, and **0C**. No PDCM array is configured for Director **B**. The PDCM array configured for Director **A** prohibits the source server from transmitting or receiving data across ISL **2**. However, internal route tables at both directors indicate a valid server-to-destination device path across ISL **1**.

A problem arises when the source server transmits Class 3 Fibre Channel data to devices across ISL **1**, consuming the ISL bandwidth. Destination devices are unaware of the PDCM array configured at Director **A** and transmit frames back to the server across ISL **2**. Because the server is prohibited from communicating across this ISL, Class 3 Fibre Channel frames are discarded without generating a busy (BSY) frame, reject (RJT) frame, or otherwise notifying the destination devices. The server receives no response from destination devices and times out. Thus, a server or device failure is indicated, when in fact the problem is a user-defined prohibited connection.

Preferred Path

The preferred path option allows a user to specify and configure one or more ISL data paths between multiple directors or fabric switches in a fabric. At each fabric element, a preferred path consists of a source port on the director or switch being configured, an exit port on the director or switch, and the Domain_ID of the destination director or switch. Each participating director or switch must be configured as part of a desired path. The following rules apply when configuring a preferred path:

- The switch Domain_ID must be set to *Insistent*.
- Domain_IDs range between **1** through **31**.
- Source and exit port numbers are limited to the range of ports available on the director or switch.
- For each source port, only one path is defined to each destination Domain_ID.

Refer to the three-director preferred path illustrated in Figure 5-6.

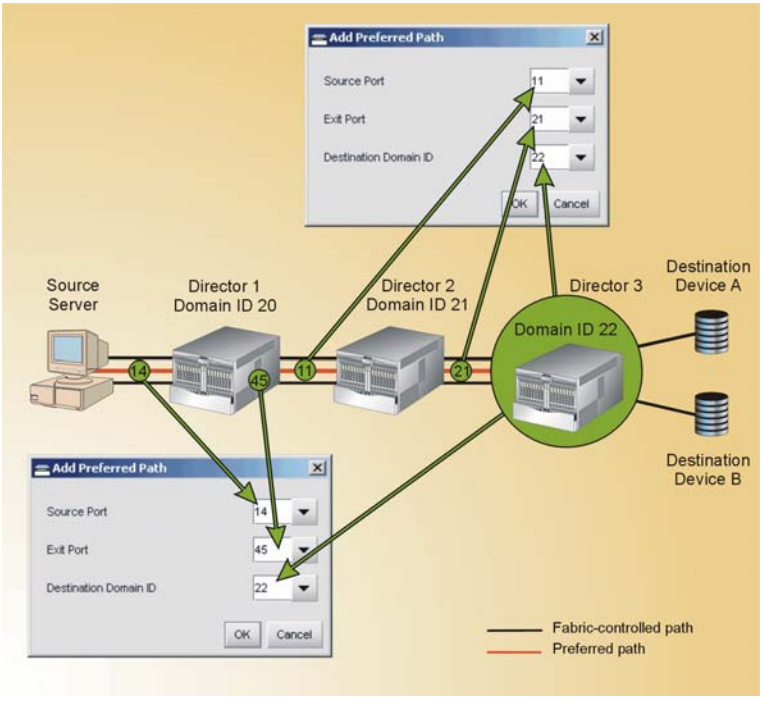

Figure 5-6 Preferred Path Configuration

A preferred path is configured between a source server and destination device (**A** or **B**), traversing Director **1**, Director **2**, and Director **3**. To configure the preferred path through the first director:

1. Select the *Preferred Path* option from the Element Manager application's *Configure* menu. The *Configure Preferred Paths* dialog box displays.
2. Click *Add*. The *Add Preferred Path* dialog box displays (bottom of Figure 5-6).
3. For the director entry port, type **14** in the *Source Port* field. For the director exit port, type **45** in the *Exit Port* field. For the destination device (Director **3**), type **22** in the *Destination Domain_ID* field.
4. Click *OK* to save the path configuration and close the dialog box.

This procedure only specifies that data enters and exits Director **1** through specific ports on the path to Director **3**. The procedure must be repeated at the second director as follows:

1. Select the *Preferred Path* option from the Element Manager application's *Configure* menu. The *Configure Preferred Paths* dialog box displays.
2. Click *Add*. The *Add Preferred Path* dialog box displays (top of Figure 5-6).
3. For the director entry port, type **11** in the *Source Port* field. For the director exit port, type **21** in the *Exit Port* field. For the destination device (Director **3**), type **22** in the *Destination Domain_ID* field.
4. Click *OK* to save the path configuration and close the dialog box.

Activating a preferred path can result in receipt of out-of-order frames (especially in FICON environments) if the path differs from the current path, if input and output (I/O) is active from the source port, and if congestion is present on the path.

To avoid problems in FICON environments, vary associated channel path identifiers (CHPIDs) temporarily offline, configure the preferred path, and vary the CHPIDs back online.

Zoning

Directors and fabric switches support a user configuration that partitions attached devices into restricted-access groups called zones. Devices in the same zone can recognize and communicate with each other through switched port-to-port connections. Devices in separate zones cannot recognize name server or route table information and therefore cannot communicate with each other. Figure 5-7 illustrates an Intrepid 6064 Director with three zones (four devices per zone).

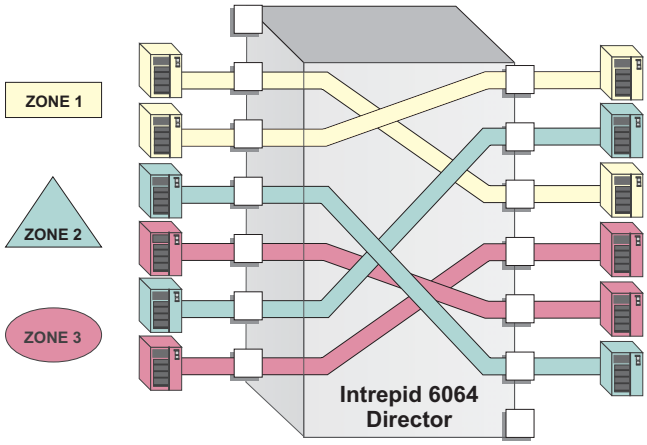

Figure 5-7 Director Zoning

Zoning is enabled and enforced by one of the following processes:

- **Software-enforced zoning -** For earlier versions of director or fabric switch firmware (prior to E/OSc Version 6.0), device configuration at a fabric element enforces zoning by limiting access to name server information in response to a device query. Only devices in the same zone as the requesting device are returned in the query response. This type of zoning is also called name server zoning or soft zoning.
- **Hardware-enforced zoning -** For later versions of director or fabric switch firmware (E/OSc Version 6.0 and later), device configuration at a fabric element enforces zoning by programming route tables that strictly prevent Fibre Channel traffic between devices that are not in the same zone. This type of zoning is also called hard zoning.

Zones are configured through the SAN management application by authorizing or restricting access to name server or route table information (depending on the firmware release level) associated with device N_Ports that attach to director or switch fabric ports (F_Ports).

Benefits of Zoning

System administrators create zones to increase network security measures, differentiate between operating systems, and prevent data loss or corruption by controlling access between devices (such as servers and data storage units), or between separate user groups (such as engineering or human resources). Zoning allows an administrator to establish:

- Logical subsets of closed user groups. Administrators can authorize access rights to specific zones for specific user groups, thereby protecting confidential data from unauthorized access.
- Barriers between devices that use different operating systems. For example, it is often critical to separate servers and storage devices with different operating systems because accidental transfer of information from one to another can delete or corrupt data. Zoning prevents this by grouping devices that use the same operating systems into zones.
- Groups of devices that are separate from devices in the rest of a fabric. Zoning allows certain processes (such as maintenance or testing) to be performed on devices in one group without interrupting devices in other groups.
- Temporary access between devices for specific purposes. Administrators can remove zoning restrictions temporarily (for example, to perform nightly data backup), then restore zoning restrictions to perform normal processes.

Configuring Zones

Zoning is configured through the SAN management application by authorizing or restricting access to name server or route table information associated with device N_Ports that attach to director or switch F_Ports or fabric loop ports (FL_Ports). A device N_Port or node loop port (NL_Port) can belong to multiple zones. Zoning is configured by:

- The eight-byte (64-digit) WWN assigned to the HBA or Fibre Channel interface installed in the device connected to the director or fabric switch.

ATTENTION ! If zoning is implemented by WWN, removal and replacement of a device HBA or Fibre Channel interface (thereby changing the device WWN) disrupts zone operation and may incorrectly exclude a device from a zone.

- The domain identification (Domain_ID) and physical port number of the director or fabric switch port to which the device is attached.

ATTENTION ! If zoning is implemented by port number, a change to the director or fabric switch fiber-optic cable configuration disrupts zone operation and may incorrectly include or exclude a device from a zone.

A zone contains a set of attached devices that can access each other. Zones are grouped into zone sets. A zone set is a group of zones that is enabled (activated) or disabled across all directors and switches in a multiswitch fabric. Only one zone set can be enabled at one time. Zone members are defined and zones or zone sets are created using the SAN management application. McDATA products support the following zoning features:

- **Zone members -** the maximum number of members configurable for a zone is 4,096.
- **Number of zones -** the maximum number of configurable zones in a zone set is 1,023 (1,024 including the default zone).
- **Number of zone sets -** the maximum number of configurable zones sets in the zoning library is 64.
- **Active zone set -** the zone set that is active across all directors and switches in a multiswitch fabric. For the active zone set:
  - — When a specific zone set is activated, that zone set replaces the active zone set.
  - — If the active zone set is disabled, all devices attached to the fabric become members of the default zone.
  - — All devices not included as members of the active zone set are included in the default zone.
- **Default zone -** the default zone consists of all devices not configured as members of a zone in the active zone set. If there is no active zone set, then all devices attached to the fabric are in the default zone. For the default zone:

  - The default zone is enabled or disabled separately from the active zone set.
  - If the default zone is enabled, then all devices not in a specified zone are included in the default zone and can communicate with each other.
  - If the default zone is disabled and there is no active zone set, then the zoning feature is completely disabled for the fabric and no devices can communicate with each other.
  - All devices are considered to be in the default zone if there is no active zone set.

- **RSCN service requests -** registered state change notification (RSCN) service requests are transmitted to all N_Ports or NL_Ports attached to the director or switch when the zoning configuration is changed.
- **Broadcast frames -** Class 3 broadcast frames are transmitted to all N_Ports attached to the director or switch, regardless of zone membership.

Joining Zoned Fabrics

Directors and fabric switches are linked through ISLs to form multiswitch fabrics. In a multiswitch fabric, the active zoning configuration applies to the entire fabric. Any change to the configuration applies to all directors and switches in the fabric.

When fabrics attempt to join, participating fabric elements exchange active zone configurations and determine if their configurations are compatible. If the configurations are compatible, the fabrics join. The resulting configuration is a single zone set containing zone definitions from each fabric. If the configurations cannot merge, E_Ports that form the ISL for each fabric element become segmented. The ports cannot transmit data frames between attached switches (class 2 or 3 traffic) but can transmit control frames (class F traffic).

Zoning configurations are compatible if there are no duplicate Domain_IDs, the active zone set name is the same for each fabric (or switch in the fabric), and zones with the same names in each fabric have identical members.

Factors to Consider When Implementing Zoning

Consider the following factors when planning to implement zoning for one or more directors or switches in the enterprise. In particular, consider the implications of zoning within a multiswitch fabric.

- **Reasons for zone implementation -** Determine if zoning is to be implemented for the enterprise. If so, evaluate if the purpose of zoning is to differentiate between operating systems, data sets, user groups, devices, processes, or some combination thereof. Plan the use of zone members, zones, and zone sets accordingly.
- **Zone members specified by port number or WWN -** Determine if zoning is to be implemented by port number or WWN. Because changes to port connections or fiber-optic cable configurations disrupt zone operation and may incorrectly include or exclude a device from a zone, zoning by WWN is recommended. However, if zoning is implemented by WWN, removal and replacement of a device HBA or Fibre Channel interface disrupts zone operation and will exclude a new device from a zone unless the device is added to the zone set.
- **Zoning implications for a multiswitch fabric -** For a multiswitch fabric, zoning is configured on a fabric-wide basis, and any change to the zoning configuration is applied to all switches in the fabric. To ensure zoning is consistent across a fabric, there can be no duplicate Domain_IDs, the active zone set name must be consistent, and zones with the same name must have identical elements. Ensure these rules are enforced when planning zones and zone sets, and carefully coordinate the zoning and multiswitch fabric tasks.

Obtaining Professional Services

Planning and implementing the zoning feature is a complex and difficult task, especially for multiswitch fabrics. Obtain planning assistance from McDATA's professional services organization before implementing the director or switch zoning feature.

Server and Storage-Level Access Control

To enhance the access barriers and network security provided by zoning through the director or fabric switch, security measures for SANs can also be implemented at servers and storage devices.

Server-level access control is called persistent binding. Persistent binding uses configuration information stored on the server and is implemented through the server's HBA driver. The process binds a server device name to a specific Fibre Channel storage volume or logical unit number (LUN), through a specific HBA and storage port WWN. For persistent binding:

- Each server HBA is explicitly bound to a storage volume or LUN, and access is explicitly authorized (access is blocked by default).
- The process is compatible with OSI standards. The following are transparently supported:
  - — Different operating systems and applications.
  - — Different storage volume managers and file systems.
  - — Different fabric devices, including disk drives, tape drives, and tape libraries.
- If the server is rebooted, the server-to-storage connection is automatically re-established.
- The connection is bound to a storage port WWN. If the fiber-optic cable is disconnected from the storage port, the server-to-storage connection is automatically re-established when the port cable is reconnected. The connection is automatically re-established if the storage port is cabled through a different director or switch port.

Access control can also be implemented at the storage device as an addition or enhancement to redundant array of independent disks (RAID) controller software. Data access is controlled within the storage device, and server HBA access to each LUN is explicitly limited (access is blocked by default). Storage-level access control:

- Provides control at the storage port and LUN level and does not require configuration at the server.
- Supports a heterogeneous server environment and multiple server paths to the storage device.
- Is typically proprietary and protects only a specific vendor's storage devices. Storage-level access control may not be available for many legacy devices.

Security Best Practices

When implementing a enterprise data security policy, establish a set of best practice conventions using methods described in this section in the following order of precedence (most restrictive listed first):

1. **SANtegrity Authentication -** The SANtegrity Authentication feature is recommended for high-security SANs to provide user-configurable, software-enforced password protection and encrypted authentication for the management server, directors, and fabric switches. These features significantly restrict access to Fibre Channel fabric elements.

2. **SANtegrity Binding -** The SANtegrity Binding feature is recommended for large and complex SANs with fabrics and devices provided by multiple OEMs or that intermix FCP and FICON protocols. The feature is required for FICON-cascaded high-integrity SANs. SANtegrity Binding includes:

   — Fabric binding (configured and enabled through the SAN management application) that allows only user-specified directors or switches to attach to specified fabrics in a SAN.

   — Switch binding (configured and enabled through the Element Manager application) that allows only user-specified devices and fabric elements to connect to specified director or fabric switch ports.

   SANtegrity Binding explicitly prohibits connections that are not user configured (unauthorized ISLs or device connections *do not* initialize and devices *do not* log in), and takes precedence over allowed connectivity in PDCM arrays, allowed connectivity through hard or soft zoning, preferred path configurations, or device-level access control.

3. **PDCM arrays -** In FICON environments, connectivity control is configured and managed at the director or fabric switch level using a PDCM array, where a user specifies which logical port addresses are allowed or prohibited from connecting with each other, including E_Port connectivity.

   Port-to-port connectivity is hardware enforced at each fabric element, and explicitly prohibited connections take precedence over allowed connectivity through hard or soft zoning, preferred path configurations, or device-level access control. However, a connection allowed through a PDCM array may be prohibited through SANtegrity Binding.

4. **Hardware-enforced zoning -** The function of hard zoning is to ensure route tables are programmed at each fabric element that explicitly allow devices to communicate *only* if the devices are in the same zone. Zoning configurations are hardware-enforced at each fabric element source port. Hard zoning impacts devices only and does not prohibit E_Port (ISL) connectivity.

   Devices in common zones can be prohibited from communicating through SANtegrity Binding or PDCM arrays, but hard zoning takes precedence over preferred path configurations, allowed connectivity through soft zoning, or device-level access control.

5. **Preferred path -** A preferred path provides soft control of fabric routing decisions on a switch-by-switch or port-by-port basis. The path instructs a fabric to use a preferred exit port out of a director or fabric switch for a specified receive port and target domain.

   If a preferred path is prohibited by SANtegrity Binding, PDCM arrays, or hard zoning, the path is not programmed. In addition, if a preferred path is not a shortest path as calculated by Dijkstra's fibre shortest path first (FSPF) algorithm, the preferred path is not programmed. However, preferred paths do take precedence over dynamic load balancing enabled through the OpenTrunking feature, soft zoning, or device-level access control.

   In general, preferred paths should be configured to influence predictable or well-known Fibre Channel traffic patterns for load balancing or distance extension applications.

6. **Software-enforced zoning -** When a device queries the name server of a fabric element for a list of other attached devices, soft zoning ensures only a list of devices in the same zone as the requesting device is returned. Soft zoning only informs a device about authorized zoning configurations; it does not explicitly prohibit an unauthorized connection. Connectivity configured through SANtegrity Binding, PDCM arrays, hardware-enforced zoning, and preferred paths takes precedence over soft zoning.

7. **Device-level access control -** Persistent binding and storage access control can be implemented at the device level as an addition or enhancement to other security features (SANtegrity Binding, PDCM arrays, zoning, and preferred paths) that are more explicitly enforced.

Security methods described in this section work in parallel with each other and are allowed to be simultaneously enabled and activated. Users are responsible for security configuration and operation within the constraints and interactions imposed by their fabric design and the methods described here.

Because incompatible security configurations can cause unintended connectivity problems or shut down Fibre Channel traffic in a fabric, it is imperative that users study and understand the interactions between SANtegrity Authentication and Binding, PDCM arrays, zoning, preferred paths, and device-level access control.

Follow best practices listed here in order of precedence. Logically work in sequence from the most restrictive method to the least restrictive method, ensuring the most restrictive connectivity or routing paths override all other paths.

Optional Feature Keys

McDATA offers several operating features that are available for the switch as customer-specified options. Available PFE keys include:

- **OSMS or FMS -** Inband director or fabric switch management is provided through purchase of the OSMS or FICON management server (FMS) feature.

  NOTE: Sphereon 4300, 4400, and 4500 Fabric Switches and SAN routers do not support out-of-band management through FMS.

- **Flexport Technology -** A Flexport Technology switch is delivered at a discount without all the ports enabled. When additional port capacity is required, the remaining ports are enabled (in four or eight-port increments) through purchase of this feature.

  NOTE: Directors and SAN routers do not support Flexport Technology.

- **SANtegrity Authentication -** Enablement of this feature enhances security in SANs by restricting unregulated access to Fibre Channel directors and fabric switches.

  NOTE: SAN routers do not support SANtegrity Authentication.

- **SANtegrity Binding -** Enablement of this feature enhances security in SANs that contain a large and mixed group of fabrics and attached devices.

  NOTE: SAN routers do not support SANtegrity Binding.

- **OpenTrunking -** Enablement of this feature provides dynamic load balancing of Fibre Channel traffic across multiple ISLs.

  NOTE: SAN routers do not support OpenTrunking.

- **Full volatility -** Enablement of this feature ensures that no Fibre Channel frames are stored after a director or Fabric switch is powered off or fails, and a memory dump file (that possibly includes classified frames) is not included as part of the data collection procedure.

  NOTE: The Intrepid 10000 Director and SAN routers do not support full volatility.

- **Full fabric -** This feature is provided only for the Sphereon 4300 Fabric Switch. Enablement of the feature provides E_Port functionality and additional port BB_Credits.
- **Remote fabric -** This feature is provided only for the Intrepid 10000 Director. Enablement of the feature provides an increased BB_Credit buffer pool and additional port credits.
- **N_Port ID virtualization -** Enablement of the N_Port ID virtualization (NPIV) feature allows up to 256 Fibre Channel addresses to be assigned to an N_Port.

  NOTE: The Intrepid 10000 Director and SAN routers do not support NPIV.

- **Element Manager application -** This feature enables director or switch management through the Element Manager user interface. Directors and switches are delivered with the application enabled for a 31-day grace period. Before grace period expiration, the application must be reactivated through a PFE key.

  NOTE: The Sphereon 4300 Fabric Switch does not support an Element Manager application.

  An Element Manager application is included with a SAN router. A PFE key is not required to enable the application.

After purchasing a feature, obtain the required PFE key through your McDATA marketing representative. A PFE key is encoded to work with the serial number of a unique director or fabric switch and is an alphanumeric string consisting of both uppercase and lowercase characters. The total number of characters may vary. The PFE key is case sensitive and must be entered exactly, including dashes.

Inband Management Access

Inband management console access (through a Fibre Channel port) is provided by enabling user-specified features that allow OSMS or FICON (FMS) host control of a director or fabric switch. The features can be simultaneously installed and enabled.

OSMS

When the OSMS feature key is enabled at the Element Manager application, host control and management of the director or switch is provided through an open-systems interconnection (OSI) device attached to a product port. When implementing inband product management through an OSI connection, plan for the following minimum host requirements:

- Connectivity to an OSI server with a product-compatible host bus adapter (HBA) that communicates through the Fibre Channel common transport (FC-CT) protocol.
- Installation of a storage network management application on the OSI server. Management applications include Veritas® SANPoint™ Control (Version 1.0 or later), or Tivoli® NetView® (Version 6.0 or later).

For information about product-compatible HBAs, third-party SAN management applications, and minimum OSI server specifications, refer to the McDATA website at www.mcdata.com.

FMS

When the FMS feature key is enabled at the Element Manager application, host control and management of the director or switch is provided through a server attached to a product port. The server communicates with the product through a FICON channel. When implementing inband product management through a FICON channel, plan for the following minimum host requirements:

- Connectivity to an IBM S/390 Parallel Enterprise Server (Generation 5 or Generation 6), with one or more FICON channel adapter cards installed, using System Automation for Operating System/390 (SA OS/390) for native FICON, Version 1.3 or later, plus service listed in the appropriate preventive service planning (PSP) bucket. The PSP bucket upgrade is HKYSA30.

  The minimum OS/390 level for a director or switch without the control unit port (CUP) feature is Version 2.6, plus service listed in PSP bucket upgrade 2032, device subset 2032OS390G5+. The minimum OS/390 level for a director or switch with the CUP feature is Version 2.1, plus service listed in the preceding PSP bucket for that function.

- Connectivity to an IBM eServer zSeries 800 (z800), zSeries 900 (z900), or zSeries 990 (z990) processor, with one or more FICON or FICON Express channel adapter cards installed, using the z/OS operating system, Version 1.1 or later.
- A host-attached Hardware Management Console. The console runs the Hardware Management Console application (HWMCA) and is the operations and management PC platform for S/390 or zSeries servers.

Flexport Technology

Sphereon 3232, 4300, and 4500 Fabric Switches can be purchased at a discount without all Fibre Channel ports enabled. The Flexport Technology feature is a hardware port expansion kit that allows customers to upgrade switch capacity on demand in eight-port increments. Flexport Technology kits are available to upgrade the:

- Sphereon 3232 Fabric Switch from 16 to 24 ports or from 24 to 32 ports.
- Sphereon 4300 Fabric Switch from four to eight ports or from eight to 12 ports.
- Sphereon 4500 Fabric Switch from eight to 16 ports or from 16 to 24 ports.

Each port expansion kit includes four or eight SFP optical transceivers, upgrade instructions, and a feature key that enables the added port capacity through the Element Manager application.

SANtegrity Authentication

SANtegrity Authentication is a feature that significantly enhances and extends SAN data security by providing password safety; CHAP or DHCHAP verification for fabric elements, management servers, and devices; a PCP user database; CT authentication for the OSMS interface; RADIUS server support; inband and out-of-band access controls lists; encrypted SSH protocol; and security logging. For additional information about the feature, refer to *SANtegrity Authentication*.

SANtegrity Binding

SANtegrity Binding is a feature that significantly enhances SAN data security. The feature includes:

- **Fabric binding -** This portion of the feature allows only specified directors or fabric switches to attach to specified fabrics in a SAN.

- **Switch binding -** This portion of the feature allows only specified devices and fabric elements to connect to specified director or fabric switch ports.
- **Enterprise Fabric Mode -** Although *Enterprise Fabric Mode* is not a keyed feature, it is required for SANtegrity Binding operation. *Enterprise Fabric Mode* also enables the following parameters:
  - — Rerouting delay.
  - — Domain RSCNs.
  - — Insistent Domain_ID.

For additional information about the feature, refer to *SANtegrity Binding*.

OpenTrunking

OpenTrunking is a feature that optimizes ISL bandwidth use in a fabric environment. The feature monitors Fibre Channel data rates (congestion and BB_Credit starvation) through multiple ISLs, dynamically applies a Dijkstra FSPF networking algorithm to calculate the optimum path between fabric elements, and load balances Fibre Channel traffic (from congested links to uncongested links) accordingly. OpenTrunking is shown in Figure 5-8.

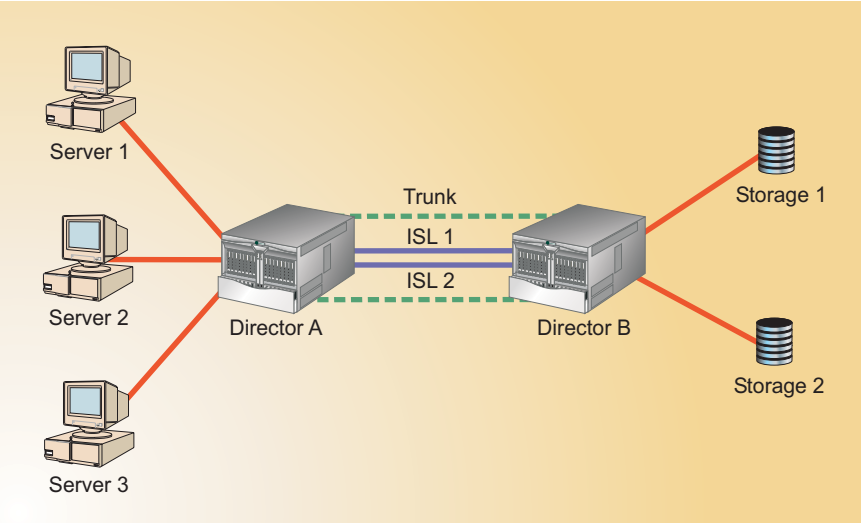

Figure 5-8 OpenTrunking

The figure illustrates two Intrepid 6064 Directors connected by two ISLs. Three servers use the ISLs to communicate with two storage devices. Without trunking, servers **1** through **3** route Fibre Channel traffic from to director B without regard to any data rates. A possible scenario is that servers **1** and **2** route high data rate traffic through ISL **1** to storage device **1** (ISL oversubscription) and server **3** routes low data rate traffic through ISL **2** to storage device **2** (ISL undersubscription).

Preferred path configurations are more restrictive than, and take precedence over, OpenTrunking. Even if OpenTrunking is enabled, no attempt is made to reroute traffic away from a preferred path, even if the path is congested or BB_Credit starved.

Full Volatility

Full volatility is a feature (available on directors and fabric switches with E/OSc Version 6.0 and later) that supports military, classified, or other high-security environments that require Fibre Channel data not be retained by the director or fabric switch after power off or failure.

When a director or fabric switch (without the full volatility feature installed) powers off or fails, a dump file is written to non-volatile random-access memory (NV-RAM). This dump file retains the last 30 Fibre Channel frames transmitted from the embedded port and the last four frames transmitted to the embedded port.

These Fibre Channel frames are then written to diskette and included as part of the data collection procedure. This process constitutes a security breach if the frame data includes classified information.

With the full volatility feature installed and enabled, no frame data is stored and the NV-RAM dump does not occur when the director or switch powers off or fails. Although this feature limits the diagnostic information available for fault isolation and resolution, the majority of failures are resolved without the dump file.

Full Fabric

The Sphereon 4300 Fabric Switch is delivered without E_Port (ISL) functionality and with all ports set to a BB_Credit value of **5**. With the full fabric feature installed and enabled, switch E_Port functionality is provided and all port BB_Credit values are increased to **12**. This feature is provided only for the Sphereon 4300 Fabric Switch.

Remote Fabric

Intrepid 10000 Director LIMs contain two scalable packet processors, each supporting an optical paddle pair. Each paddle pair provides 16 ports (1.0625 or 2.12500 Gbps operation), four ports (10.2000 Gbps operation), or ten ports (mixed data rate operation). A minimal BB_Credit buffer pool is allocated among all paddle-pair ports that allows a 1.0625 or 2.12500 Gbps port to be set to **60** BB_Credits and a 10.2000 Gbps port to be set to **360** BB_Credits.

With the remote fabric feature installed and enabled, the buffer pool is increased and each paddle pair is allocated a maximum of **1,373** BB_Credits. Feature enablement allows a long-link 1.0625 or 2.12500 Gbps port to be set to **1133** BB_Credits and a long-link 10.2000 Gbps port to be set to **1085** BB_Credits. This feature is provided only for the Intrepid 10000 Director.

N_Port ID Virtualization

NPIV is a feature that allows a single physical port to support up to 256 virtual Fibre Channel addresses. A virtual node port (NV_Port) retains full N_Port capabilities and can register for full fabric services. However, simultaneous loop device operation (NL_Port) and NPIV operation (NV_Port) is not allowed on the same physical port. This feature is beneficial in mainframe environments where multiple partitions share the same physical channel and port connection.

Element Manager Application

The Element Manager feature allows director or fabric switch management through an Element Manager application GUI. A director or switch is delivered with the application enabled for a 31-day grace period. Before grace period expiration, the application must be reactivated and enabled through a PFE key.

During the grace period, a *No Feature Key* dialog box appears when the Element Manager application is accessed. Click *OK* to close the dialog box and open the application.

In addition, the message **Element Manager license key has not been installed - Please follow up instructions to update permanent key** appears splashed across views, indicating the Element Manager PFE key must be installed.

6 Configuration Planning Tasks

This chapter describes configuration planning best-practices tasks to be performed before installing one or more McDATA Fibre Channel switching products in a storage area network (SAN) configuration. Table 6-1 summarizes planning tasks described in the chapter.

Table 6-1 Configuration Planning Tasks

Task	Page
Task 1: Prepare a Site Plan	6-2
Task 2: Plan Fibre Channel Cable Routing	6-3
Task 3: Consider Interoperability with Fabric Elements and End Devices	6-4
Task 4: Plan Console Management Support	6-5
Task 5: Plan Ethernet Access	6-7
Task 6: Plan Network Addresses	6-7
Task 7: Plan SNMP Support (Optional)	6-10
Task 8: Plan E-Mail Notification (Optional)	6-11
Task 9: Establish Product and Server Security Measures	6-11
Task 10: Plan Phone Connections	6-12
Task 11: Diagram the Planned Configuration	6-12
Task 12: Assign Port Names and Nicknames	6-13

Table 6-1 Configuration Planning Tasks (*continued*)

Task	Page
Task 13: Complete the Planning Worksheet	6-14
Task 14: Plan AC Power	6-28
Task 15: Plan a Multiswitch Fabric (Optional)	6-29
Task 16: Plan Zone Sets for Multiple Products (Optional)	6-30
Task 17: Plan SAN Routing (Optional)	6-31
Task 18: Complete Planning Checklists	6-34

Task 1: Prepare a Site Plan

For each director, fabric switch, SAN router, or FC-512 Fabricenter equipment cabinet installed, design a site plan that provides efficient work flow, operator convenience and safety, and adequate service clearances for the equipment cabinet. A customer manager should review the site plan with a service representative and consider:

- Location and relationship of the physical facilities such as walls, doors, windows, partitions, furniture, and telephones.
- Proximity of the director, fabric switch, or SAN router to servers and storage peripherals, and if a multiswitch fabric is to be enabled, proximity of participating fabric elements to each other.
- Location of at least one analog phone line (capable of providing long-distance service) for the management server to support the call-home feature or provide remote dial-in support. In addition, consider accessibility to a second phone to aid in installation and serviceability.
- Availability of Ethernet local area network (LAN) connections and cabling to support remote user workstation and simple network management protocol (SNMP) management station access. Remote user and SNMP workstations are optional.
- Equipment cabinet locations, Ethernet cabling, and the Internet protocol (IP) addressing scheme to support optional cabinet interconnection and management server consolidation.

- Power requirements, including an optional uninterruptable power supply (UPS).
- Lengths of power cables and location of electrical outlets (for directors, switches, and the management server) having the proper phase, voltage, amperage, and ground connection.

DANGER

Use the supplied power cords. Ensure the facility power receptacle is the correct type, supplies the required voltage, and is properly grounded.

- Security necessary to protect the installation's physical integrity, while maintaining accessibility to the director or switch.
- Facility access and security clearances for installation personnel.
- Equipment cabinet front and rear service clearances, operator clearances, and maintenance access clearances.
- Weight of a Fabricenter equipment cabinet. Either multiple persons or a lift must be available during installation to remove the cabinet from the packing crate.
- Heat dissipation, temperature, and humidity requirements.

For specific actions required to satisfy this planning task, refer to Table 6-2 (Physical Planning and Hardware Installation Tasks), and Table 6-3 (Operational Setup Tasks).

Task 2: Plan Fibre Channel Cable Routing

Plan for sufficient singlemode fiber-optic, multimode fiber-optic, and Ethernet cabling to meet the connectivity requirements for all fabric elements (directors, fabric switches, and SAN routers), Fibre Channel servers, and devices. Plan for sufficient fiber-optic cabling to meet interswitch link (ISL) connectivity requirements and SAN routing requirements.

Plan for at least one meter (39.37 inches) of fiber-optic cable inside the Fabricenter equipment cabinet for routing to product Fibre Channel ports as required. Plan for an additional 1.5 meters (5 feet) of cable outside the cabinet to provide slack for service clearance, limited cabinet movement, and inadvertent cable pulls.

In addition, consider the following when planning cable routing:

- The need for additional fiber-optic cables could grow rapidly. Consider installing cable with extra fibers, especially in hard to reach places like underground trenches. Consider locating the equipment cabinet near a fiber-optic patch panel.
- Follow proper procedures when moving an installed equipment cabinet to prevent cable or connector damage.

Task 3: Consider Interoperability with Fabric Elements and End Devices

McDATA conducts a substantial level of testing to ensure director, fabric switch, and SAN router interoperability with fabric elements and end devices provided by multiple original equipment manufacturers (OEMs). New devices are tested and qualified on a continual basis. Contact your McDATA representative for the latest information about fabric element, server, host bus adapter (HBA), and device interoperability.

Consider whether to set the director or fabric switch to open systems or fibre connection (FICON) management style. This setting only affects the operating style used to manage the product; it does not affect port operation. Open-systems interconnection (OSI) devices can communicate with each other if the product is set to FICON management style, and FICON devices can communicate with each other if the product is set to open systems management style. Be aware that:

- When a director or switch is set to open systems management style, a traditional Fibre Channel fabric consisting of multiple domains (fabric elements) is supported. Inband management through the open-systems management server (OSMS) is also supported.
- When a director or switch is set to FICON management style, only a single domain (fabric element) is supported. Inband management through the FICON management server (FMS) is also supported. When operating using FICON management style, ports are set to F_Port operation, thus eliminating E_Port, ISL, and multiswitch fabric capabilities.

NOTE: If the FICON management server feature is enabled, the default operating style is FICON. The open systems management style cannot be enabled.

Consider purchasing and enabling the SANtegrity Authentication and SANtegrity Binding features to provide additional data security in a complex and multi-OEM environment. Contact your McDATA representative for information about the features.

Task 4: Plan Console Management Support

Plan to implement one or more of the following to provide console management and support for directors, switches, and SAN routers:

- **Management Server -** The 1U rack-mount management server is used for product installation, initial software configuration, changing the configuration, and monitoring product operation.
  - — When SAN management and Element Manager applications (directors and fabric switches) or the SANvergence Manager application (SAN routers) are installed on the management server, the server is used as a local user workstation.
  - — The management server can support up to 48 managed McDATA products.
  - — Managed products can be powered off and on without the management server.
  - — A management server failure does not cause an operating director, switch, or SAN router to fail.
  - — The management server is fully operational, even if there is no user logged in to the Windows 2000 Professional operating system. The server allows remote users to log in and continues to monitor products in the background.

NOTE: The Sphereon 4300 Fabric Switch is not supported by the management server.

  - — Ensure SAN router-specific requirements for management server support are evaluated. Refer to *Task 17: Plan SAN Routing (Optional)* for information.

- **Remote user workstations -** If remote access to the management server is required, plan to install user workstations with the SAN management and Element Manager applications configured. Administrators can use these remote workstations to configure and monitor directors and fabric switches. Up to 25 sessions can be simultaneously active. Sessions from remote user workstations are disabled if the management server is powered off.

  NOTE: Remote management server access to SAN routers is not supported.

- **Inband management support -** If inband console management of a director or fabric switch is required, plan for a Fibre Channel port connection that communicates with the attached server.

  If director or fabric switch management through an OSI server is planned, ensure the OSMS feature key is ordered with the Element Manager application. This feature enables host control of the product from an OSI server attached to a Fibre Channel port. Ensure the server meets minimum specifications and a product-compatible HBA and appropriate operating system or SAN management application is available.

  If director or switch management through an IBM host is planned, ensure the FMS feature key is ordered with the Element Manager application. This feature key enables host control of the product from an IBM System / 390 Parallel Enterprise Server or eServer zSeries processor attached to a Fibre Channel port.

- **EFCM Basic Edition interface -** If a web browser-capable PC and Internet access to a product EFCM Basic Edition interface are required, plan accordingly and ensure access to an analog phone line. Access to the EFCM Basic Edition interface is not provided by the management server.

  NOTE: EFCM Basic Edition interface access to the Intrepid 10000 Director and SAN routers is not supported.

Task 5: Plan Ethernet Access

The management server and one or more products are configured on a dedicated Ethernet LAN segment and delivered in the Fabricenter equipment cabinet. No Ethernet access planning is required for a stand-alone cabinet. This task is required to:

- **Connect equipment cabinets -** Ethernet hubs in multiple equipment cabinets can be connected to provide management server access to up to 48 managed McDATA products. Cabinets can be placed at any distance up to the limit of the 10/100 megabit per second (Mbps) LAN segment.
- **Consolidate management server operation -** If management server operation is to be consolidated to one primary server and one or more backup servers, plan for Ethernet cabling to interconnect equipment cabinets and ensure all directors, switches, and server platforms have unique IP addresses.
- **Install equipment cabinets on a public LAN -** If a public LAN segment is to be used, determine from the customer's network administrator how to integrate the products and management server. Ensure all access, security, and IP addressing issues are resolved.

  NOTE: It is recommended that directors, fabric switches, SAN routers, and the management server be installed on a dedicated Ethernet hub and LAN segment to avoid security, traffic, and fault isolation problems associated with a public LAN.

- **Install remote user workstations -** Plan for access to the LAN segment (dedicated or public) containing the management server if remote user workstations are required.

Task 6: Plan Network Addresses

Depending on the configuration of the LAN on which directors, switches, SAN routers, and the management server are installed, plan network addressing as follows:

- If installing products and the management server on a dedicated (private) LAN segment, there is no requirement to change any default network addresses. However, if multiple equipment cabinets are connected, ensure all managed products and servers have unique IP addresses. If new IP addresses are required, consult with the customer's network administrator.
- If installing products and the management server on a public LAN containing other devices, default network addresses may require change to avoid address conflicts with existing devices.

  For Intrepid-series directors, Sphereon-series fabric switches, and Eclipse-series SAN routers, the IP address, gateway address, and subnet mask are changed through a remote terminal connected to the product's maintenance port.

  For the management server, these addresses are changed through the liquid crystal display (LCD) front panel. In addition, assign and record a unique domain name system (DNS) name for the management server and each director or switch.
- Gateway addresses may need to be configured for managed products and the management server if these devices connect to the LAN through a router or other gateway device.

The Ethernet connections for the 1U management server, directors, fabric switches, and SAN routers have the following default network addresses:

- **1U Management server:**
  - — Media access control (MAC) address is unique for each server and managed product. The address is in ***xx.xx.xx.xx.xx.xx*** format, where each ***xx*** is a hexadecimal pair.
  - — IP address of the private LAN connection (LAN 2) is **10.1.1.1**.
  - — IP address of the public LAN connection (LAN 1) is **192.168.0.1**.
  - — Subnet mask is **255.0.0.0**.
  - — Gateway address is blank.

- **Intrepid-series directors and Sphereon-series fabric switches:**
  - MAC address is unique for each product.
  - Default IP address is **10.1.1.10**.
  - Subnet mask is **255.0.0.0**.
  - Gateway address is **0.0.0.0**.
- **Eclipse 1620 SAN Router:**
  - **System addresses:**
    - MAC address is unique for each product.
    - Default IP address is **192.168.111.100**.
    - Subnet mask is **255.255.255.0**.
    - Gateway address is **0.0.0.0**.
  - **10/100 Base-T Ethernet management port addresses:**
    - Default IP address is **192.168.100.100**.
    - Subnet mask is **255.255.255.0**.
    - Gateway address is **0.0.0.0**.
  - **Intelligent port (3) addresses:**
    - Default IP address is **0.0.0.0**.
    - Subnet mask is **0.0.0.0**.
    - External IP address is **0.0.0.0**.
    - Internal IP address is **192.168.111.103**.
  - **Intelligent port (4) addresses:**
    - Default IP address is **0.0.0.0**.
    - Subnet mask is **0.0.0.0**.
    - External IP address is **0.0.0.0**.
    - Internal IP address is **192.168.111.104**.
- **Eclipse 2640 SAN Router:**
  - **System addresses:**
    - MAC address is unique for each product.
    - Default IP address is **0.0.0.0**.

- Subnet mask is **0.0.0.0**.
- Gateway address is **0.0.0.0**.

— **10/100 Base-T Ethernet management port addresses:**

- Default IP address is **192.168.100.100**.
- Subnet mask is **255.255.255.0**.
- Gateway address is **0.0.0.0**.

— **Intelligent port (13 through 16) addresses:**

- Default IP address is **0.0.0.0**.
- Subnet mask is **0.0.0.0**.
- External IP address is **0.0.0.0**.
- Internal IP address is **0.0.0.0**.

Task 7: Plan SNMP Support (Optional)

As an option, network administrators can use the SAN management application to configure an SNMP agent that runs on the management server. This agent is used to:

- Configure up to 12 authorized management workstations to receive unsolicited SNMP trap messages (directors and fabric switches).
- Configure up to eight authorized management workstations to send SNMP trap messages (SAN routers).

Administrators can also use the Element Manager application to configure an SNMP agent that runs on each director, fabric switch, or SAN router. This agent can be configured to send generic SNMP trap messages. Trap recipients can also access SNMP management information and may be granted permission to modify SNMP variables as follows:

- Assign and record product names, contact persons, descriptions, and locations to configure the products for SNMP management station access.
- Plan access to the managed product LAN segment. This segment must connect to the LAN on which SNMP management workstations are installed.

- Obtain IP addresses and SNMP community names for management workstations that have access to products.
- Determine which (if any) management workstations can have write permission for SNMP variables.
- Obtain product-specific trap information from McDATA to load onto SNMP management workstations.

Task 8: Plan E-Mail Notification (Optional)

As an option, network administrators can configure director and fabric switch e-mail support. The following support considerations are required if the e-mail notification feature is used:

- Determine if e-mail notification is to be configured and used for significant system events.
- Determine which persons (up to five) require e-mail notification of significant director or switch events and record their e-mail addresses.
- Identify an attached e-mail server that supports the simple mail transfer protocol (SMTP) standard as defined in RFC 821.

NOTE: E-mail notification for SAN routers is not supported.

Task 9: Establish Product and Server Security Measures

Effective network security measures are recommended for directors, fabric switches, SAN routers, and the management server. Physical access to the network should be limited and monitored, and password control should be strictly enforced. When planning security measures, consider the following:

- Managed products and the management server are installed on a LAN segment and can be accessed by attached devices (including devices connected through a remote LAN). Access from remote devices is limited by installing the management server and managed products on an isolated, dedicated LAN segment. This approach is recommended. Installing and enabling the SANtegrity Authentication and SANtegrity Binding features is also recommended.

- Remote access to products is possible through the maintenance port or an internal modem connection to the management server. These connections are for use by authorized service personnel only and should be carefully monitored.
- The number of remote workstations with access to the management server and managed products can and should be restricted. Obtain IP addresses for workstations that should have exclusive access. Ensure adequate security measures are established for the configured workstations.
- Carefully manage users (up to 16) who have access to the SAN management, SANvergence Manager, and Element Manager applications, and assign user names, passwords, and user rights.
- Ensure adequate security controls are established for remote access software, including the EFCM Basic Edition interface.

Task 10: Plan Phone Connections

Analog telephone connections are used by service personnel and for access to the management server's internal modem. Plan for one or more telephone connections near the server because of the following:

- If a field-replaceable unit (FRU) in a managed product fails, the management server provides a call-home feature that transmits a message through the server's internal modem connection to a designated support center.
- While performing a diagnostic or repair action, a service representative or network administrator at the management server may require voice technical support through a telephone connection.
- A service representative may need to connect to the management server through the internal modem to access maintenance and utility functions, check status, and perform other tasks.

Task 11: Diagram the Planned Configuration

Determine peripheral devices that will connect to each director, switch, or SAN router and where connectivity should be limited (zoning). These devices may include servers, storage control devices, and other fabric elements in a multiswitch fabric.

Part of this task may have been performed when the configuration was determined. It might be helpful to draw the configuration diagram. Indicate distances in the diagram if necessary. Transfer information from the configuration diagram to the product planning worksheet provided as part of *Task 13: Complete the Planning Worksheet*.

Task 12: Assign Port Names and Nicknames

Consider assigning names to director, switch, or SAN router ports based upon devices connected to the ports. Though not required, port naming provides convenience and ease of use. Port naming also documents devices that connect through individual ports and identifies what is attached to each port. When it is necessary to change port connectivity, port names make it easier to identify the ports and attached end devices.

Also consider assigning nicknames to device and fabric worldwide names (WWNs). Though not required, nicknaming provides a useful substitute for the cryptic eight-byte WWN. Once a nickname is assigned, it is referenced throughout the SAN management application.

Transfer port names and nicknames to the product planning worksheet provided as part of *Task 13: Complete the Planning Worksheet*.

Rules for Port Names

Port names can be up to 24 alphanumeric characters in length. Spaces, hyphens, and underscores are allowed within the name. Each port name must be unique for a product. However, the same port name can be used on separate products. It is recommended that unique port names be used, particularly within a complex multiswitch fabric. Example port names include:

Lab server.

Test system-2.

Printer_001.

Rules for Nicknames

Nicknames can be up to 32 alphanumeric characters in length. Spaces, hyphens, and underscores are allowed within the nickname. Each nickname must be unique (corresponding to a unique WWN). Example nicknames include:

Fabric-1.

Host system-1.

DASD_001.

Task 13: Complete the Planning Worksheet

The planning worksheet included in this task is an eight-page form that depicts port assignments for a director, switch, or SAN router. The worksheet lists 256 ports, equal to the capability of the highest port-count product described in this publication. The worksheet provides fields to identify devices that connect to the ports.

Transfer information from the configuration diagram (completed while performing *Task 11: Diagram the Planned Configuration*) to the worksheet, and transfer port names and nicknames (assigned while performing *Task 12: Assign Port Names and Nicknames*). In addition, indicate all unused ports. Retain the planning worksheet as part of a permanent record.

Product Planning Worksheet (Page 1 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
00						
01						
02						
03						
04						
05						
06						
07						
08						
09						
10						
11						
12						
13						
14						
15						
16						
17						
18						
19						

Product Planning Worksheet (Page 2 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
20						
21						
22						
23						
24						
25						
26						
27						
28						
29						
30						
31						
32						
33						
34						
35						
36						
37						
38						
39						

Product Planning Worksheet (Page 3 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
40						
41						
42						
43						
44						
45						
46						
47						
48						
49						
50						
51						
52						
53						
54						
55						
56						
57						
58						
59						

Product Planning Worksheet (Page 4 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
60						
61						
62						
63						
64						
65						
66						
67						
68						
69						
70						
71						
72						
73						
74						
75						
76						
77						
78						
79						

Product Planning Worksheet (Page 5 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
80						
81						
82						
83						
84						
85						
86						
87						
88						
89						
90						
91						
92						
93						
94						
95						
96						
97						
98						
99						

Product Planning Worksheet (Page 6 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
100						
101						
102						
103						
104						
105						
106						
107						
108						
109						
110						
111						
112						
113						
114						
115						
116						
117						
118						
119						

Product Planning Worksheet (Page 7 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
120						
121						
122						
123						
124						
125						
126						
127						
128						
129						
130						
131						
132						
133						
134						
135						
136						
137						
138						
139						

Product Planning Worksheet (Page 8 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
140						
141						
142						
143						
144						
145						
146						
147						
148						
149						
150						
151						
152						
153						
154						
155						
156						
157						
158						
159						

Product Planning Worksheet (Page 9 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
160						
161						
162						
163						
164						
165						
166						
167						
168						
169						
170						
171						
172						
173						
174						
175						
176						
177						
178						
179						

Product Planning Worksheet (Page 10 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
180						
181						
182						
183						
184						
185						
186						
187						
188						
189						
190						
191						
192						
193						
194						
195						
196						
197						
198						
199						

Product Planning Worksheet (Page 11 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
200						
201						
202						
203						
204						
205						
206						
207						
208						
209						
210						
211						
212						
213						
214						
215						
216						
217						
218						
219						

Product Planning Worksheet (Page 12 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____			Attached Devices			
Port	**Port Name**	**Location**	**Type**	**Model**	**IP Address**	**Zone**
220						
221						
222						
223						
224						
225						
226						
227						
228						
229						
230						
231						
232						
233						
234						
235						
236						
237						
238						
239						

Product Planning Worksheet (Page 13 of 13)

Switch Name:_____ IP Address:_____ Unit Name:_____	Attached Devices

Port	Port Name	Location	Type	Model	IP Address	Zone
240						
241						
242						
243						
244						
245						
246						
247						
248						
249						
250						
251						
252						
253						
254						
255						

Task 14: Plan AC Power

Plan for facility power sources for each Fabricenter equipment cabinet, director, fabric switch, or SAN router as follows:

- The Fabricenter equipment cabinet operates at 47 to 63 Hertz (Hz), 200 to 240 volts alternating current (VAC), and requires a minimum dedicated 30-ampere service.
- The Intrepid 6140 Director operates at 47 to 63 Hz, 200 to 240 VAC, and requires a minimum dedicated 15-ampere service.
- The Intrepid 10000 Director operates at 47 to 63 Hz, 200 to 240 VAC, and requires a minimum dedicated 20-ampere service.
- Other directors, fabric switches, and SAN routers in the cabinet operate at 47 to 63 Hz, 100 to 240 VAC, and require a minimum dedicated 5-ampere service.

If two power sources are supplied (optional but recommended for high availability), the equipment cabinet contains two customer-specified external power cords. Each cord should be connected to a separate power circuit, or both should be connected to a UPS. Several types of power cables and plugs are available to meet local electrical requirements.

DANGER

Use the supplied power cords. Ensure the facility power receptacle is the correct type, supplies the required voltage, and is properly grounded.

Keep all power cables out of high-traffic areas for safety and to avoid power interruption caused by accidentally unplugging the product or Fabricenter equipment cabinet.

Task 15: Plan a Multiswitch Fabric (Optional)

If a multiswitch fabric topology is to be implemented, carefully plan the physical characteristics and performance objectives of the topology, including the proposed number of fabric elements, characteristics of attached devices, cost, nondisruptive growth requirements, and service requirements. Refer to *Fabric Topologies*, *Planning for Multiswitch Fabric Support*, and *General Fabric Design Considerations* for detailed information.

When two or more fabric elements are connected through ISLs to form a fabric, the elements must have compatible operating parameters, compatible name server zoning configurations, and unique domain identifications (IDs). Planning for a fabric must be carefully coordinated with planning for zoned configurations. The following factors should be considered when planning for a multiswitch fabric:

- **Fabric topology limits -** Consider the practical number of fabric elements (theoretical maximum of 31, practical limit of 24), number of ISLs per element, hop count (maximum of three), and distance limitations (limited by port type and cable availability).
- **Multiple ISLs for bandwidth and load balancing -** Consider using multiple ISLs to increase the total bandwidth available between two fabric elements. If heavy traffic between devices is expected, also consider multiple ISLs to create multiple minimum-hop paths for load balancing.

  If multiple ISL connections are planned, ensure the OpenTrunking feature key is ordered with the Element Manager application. This feature automatically provides dynamic load balancing across multiple ISLs in a fabric environment.
- **Principal switch selection -** If required, plan which fabric element is to be assigned principal switch duties for the fabric.
- **Critical operations -** Consider routing paths that transfer data for critical operations directly through one director or fabric switch and not through the fabric.

Planning and implementing a multiswitch fabric is a complex and difficult task. Obtain planning assistance from McDATA's professional services organization before implementing a fabric topology.

Task 16: Plan Zone Sets for Multiple Products (Optional)

If name server zoning is to be implemented, carefully plan the characteristics and security objectives (separation of operating systems, data sets user groups, devices, or processes) of zone members, zones and zone sets.

If a fabric topology or routed SAN is implemented, zoning is configured on a fabric-wide or SAN-wide basis. Planning for zoned configurations must be carefully coordinated with planning the topology. The following factors should be considered when planning to implement name server zoning:

- **Zone members specified by port number or WWN** - Consider if zoning is to be implemented by port number or WWN. Because changes to a port connections or fiber-optic cable configurations may disrupt zone operation, zoning by WWN is recommended.

  NOTE: SAN routers do not support port number zoning.

- **Zoning implications for a multiswitch fabric** - To ensure zoning is consistent across a multiswitch fabric, directors and fabric switches must have compatible operating parameters and unique domain IDs, the active zone set name must be consistent, and zones with the same name must have identical elements.
- **Zoning implications for a routed SAN** - A zone policy must be established that specifies how zone information is synchronized between a SAN router and attached fabrics. Zone policy options are **No Zone Synchronization** (device zoning is controlled at the fabric level) or **Append IPS Zones** (device zoning control is shared between a SAN router and the fabric).
- **Server and storage device access control** - In addition to zoning, consider implementing server-level access control (persistent binding) and storage-level access control.

Consider purchasing and enabling the SANtegrity Authentication and SANtegrity Binding features to work in conjunction with name server zoning to provide additional data security in a complex and multi-OEM environment. Planning and implementing zones and zone sets is a complex and difficult task, especially for multiswitch fabrics. Obtain planning assistance from McDATA's professional services organization before implementing a zoning feature.

Task 17: Plan SAN Routing (Optional)

If a routed SAN is to be implemented, carefully plan metropolitan area network (MAN) or wide area network (WAN) connectivity and the integration of SAN routers with standard fibre channel fabric elements.

Ensure basic requirements for physical SAN routers are incorporated in the site plan(s) developed while performing *Task 1: Prepare a Site Plan*. Plan equipment cabinet locations; availability and location of Ethernet LAN connections and cabling; and power requirements. The following router-specific factors should also be considered:

- **Management server support -** At each location that requires server support for one or more SAN routers, consider the following:
  - The 10/100 Base-T Ethernet management port on each SAN router provides out-of-band connectivity to the management server. Gigabit Ethernet (GbE) intelligent ports on each SAN router provide inband IP network connectivity to transmit storage traffic. The management port and intelligent ports require network address configuration, including an IP address and subnet mask. It is recommended the management port be configured on a separate subnet from intelligent ports.
  - The SANvergence Manager application (management server resident) and Element Manager application (router resident) use SNMP and hypertext transfer protocol (HTTP) to communicate. To ensure communication, Java runtime environment **j2re-1_4_2_01-windows-i586-iftw.exe** (or later) must be installed on the management server. The software is downloaded from http://java.sun.com.
  - The management server port communicating with a SAN router must be set to **Auto Negotiate** to ensure viable server-to-router communication.
- **Fibre Channel network connections -** At each location, fiber-optic cables with appropriate connectors must be routed between Fibre Channel elements (storage devices, servers, directors, and fabric switches) and the SAN router. Eclipse-series SAN routers support small form factor pluggable (SFP) optical transceivers with LC duplex connectors.

- **IP network connections -** At each location, cables with appropriate connectors must be routed between the IP transport network and the SAN router. The Eclipse 1620 SAN Router supports Ethernet RJ-45 connectors or SFP optical transceivers with LC duplex connectors. The Eclipse 2640 SAN Router supports SFP optical transceivers with LC duplex connectors.
- **Establish operational mode and transport technology -** Establish if the operational mode is expected to support synchronous remote data replication (RDR/S) or asynchronous remote data replication (RDR/A). Based on operational mode requirement, establish the IP WAN transport technology as follows:
  - — Repeated or unrepeated dark fiber.
  - — Wavelength division multiplexing (WDM).
  - — Synchronous optical network (SONET) and synchronous digital hierarchy (SDH).
  - — Internet protocol.

  Refer to *Extended-Distance Operational Modes* and *SAN Extension Transport Technologies* for detailed information.
- **Determine peak available bandwidth -** The peak available bandwidth for data transport (exclusive of protocol overhead) must be determined or obtained from the network service provider. If the IP WAN link is dedicated, the peak available bandwidth equals the total link bandwidth. This implies that no other application data or traffic is routed across the link. If the IP WAN link is shared, the peak available bandwidth equals that portion of the total link bandwidth allotted for storage traffic at peak use time.

  Data ingress must not exceed peak available bandwidth or downstream network device buffers fill, overflow, and drop data packets. Dropped packets cause congestion and result in reduced link throughput. To prevent this problem, enable rate limiting to ensure the ingress data rate does not exceed the egress rate of the ***slowest*** link in the IP WAN path.

  Refer to *IRL Optimization* and *Intelligent Port Speed* for detailed information.

- **Negotiate SLA -** Network service providers provide IP WAN transport services in accordance with a negotiated service level agreement (SLA). Ensure the SLA specifies the link availability, peak available bandwidth, latency, security level, monitoring level, packet loss, and mean time to repair (MTTR).
- **Intelligent port addresses -** Each intelligent port that supports Internet Fibre Channel protocol (iFCP) requires an IP address and subnet mask. Depending on the gateway (next hop router), the port may require an internal IP address and external IP address.

  If there is a layer 2 connection between SAN routers (such as WDM, SONET, or SDH), no gateway addressing is required because there are no intervening layer 3 devices and iFCP ports are on the same subnet. If there is a layer 3 connection between SAN routers, gateway addressing must be specified because iFCP ports are on different subnets.
- **Configure and test transport network -** The IP WAN transport network between SAN routers must be configured, operational, tested, and able to support bidirectional storage traffic. Specifically:

  — SAN routers must be correctly configured and able to route traffic between end subnets.

  — Devices that are part of the IP WAN infrastructure should be set to the appropriate operational mode, symmetrical flow control should be enabled, and ports should be correctly designated to support storage traffic or management traffic.

  — As part of a test plan, ping the IP WAN to ensure subnet connectivity.
- **Establish zone policy -** SAN Routing provides flexibility with respect to zoning behavior and interactions between a router and attached fabrics. The zone policy specifies how zoning information is synchronized between a SAN router and attached fabrics. It is not a requirement that all router-attached fabrics use the same zone policy. Refer to *Routed SAN Zoning* for detailed information.
- **Provide accurate documentation -** Accurate and up-to-date documentation that records all facility locations, contact personnel, device names, network addresses, port numbers, link types, cable types, protocols, and equipment makes and models is required.

Task 18: Complete Planning Checklists

As a guide for planning tasks, complete the planning checklists under this task. Checklists provide detailed planning activities and provide space for a planned completion date for each activity. The customer's management information system (MIS) project manager should examine the checklists and determine the personnel and resources required for completing planning and installation tasks. Customer personnel might be used from the following functional areas:

- Systems programming personnel to update input/output (I/O) definitions to identify directors, fabric switches, and SAN routers.
- Ethernet management personnel to obtain IP addresses, gateway addresses, and subnet masks for directors, fabric switches, SAN routers, and the management server; and a DNS host name for the server.
- Facilities planning personnel to outline the facility floor plan and to arrange for electrical wiring, receptacles and telephone lines.
- Installation planning personnel to determine fiber-optic and Ethernet cabling requirements, routing requirements, and to plan connectivity between directors, fabric switches, SAN routers, and attached devices.
- Trainers to determine training and education needs for operations, administration, and maintenance personnel.
- Administrators to determine director port names and WWN nicknames, identify attached devices, and assign password levels and user names for director, fabric switch, and SAN router access.

Table 6-2 lists physical planning and hardware installation tasks and includes the task owner, due date, and comments. Table 6-3 lists operational setup tasks and includes the task owner, due date, and comments.

Table 6-2 Physical Planning and Hardware Installation Tasks

Activity	Task Owner	Due Date	Comments
Locate the physical facilities.			
Connect the facility alternating current (AC) power circuits.			If more than one managed product, consider separate power circuits for availability.
Obtain an uninterruptable power supply (optional).			Recommended.
Obtain two outside-access phone lines.			One for the modem and the second for a telephone.
Order and deliver Ethernet and fiber-optic cables with appropriate connectors. Cables must support Fibre Channel network, management network, and IP WAN network (if SAN routing is supported) connectivity.			Cables are purchased by the customer separately. Plan to have cables installed before equipment cabinet delivery.
Order the Fabricenter cabinet with one or more McDATA managed products.			
Order Fibre Channel devices and peripherals.			
Determine proximity of the equipment cabinet (with directors, fabric switches, and SAN routers) to attached devices (multimode shortwave laser or singlemode longwave laser).			500 meters (1.0625 Gbps), 300 meters (2.1250 Gbps), 150 meters (4.2500 Gbps), or 82 meters (10.2000 Gbps) for 50/125 mm multimode cable. 300 meters (1.0625 Gbps), 150 meters (2.1250 Gbps), 75 meters (4.2500 Gbps), or 33 meters (10.2000 Gbps) for 62.5/125 mm multimode cable. 10, 20, or 35 kilometers for 9/125 mm singlemode cable, depending on the optic purchased.
Install Fibre Channel devices and peripherals.			
Set up server peripheral for inband director or switch management (optional).			Order OSMS or FMS product feature enablement (PFE) key.
Route fiber-optic jumper cables.			
Set up local area network (LAN) connections for directors, fabric switches, SAN routers, and the management server.			If SAN routing is supported, provide additional management server configuration tasks.

Table 6-2 Physical Planning and Hardware Installation Tasks (*continued*)

Activity	Task Owner	Due Date	Comments
Set up LAN connections to corporate intranet for remote workstation access (optional).			Remote workstation access is supported for directors and fabric switches only.
Determine peak available bandwidth of the available IP WAN network (optional).			If SAN routing is supported, rate limiting must be configured and enabled based on peak available bandwidth.
Negotiate SLA (optional).			If SAN routing is supported, an SLA must be negotiated with a network service provider to ensure reliable IP WAN transport service.

Table 6-3 Operational Setup Tasks

Activity	Task Owner	Due Date	Comments
Obtain or assign IP addresses, subnet masks, and gateway addresses for products.			Management server (if installing on a LAN with non-McDATA devices). Directors, fabric switches, and SAN routers (if installing on a LAN with non-McDATA devices). Remote user workstation (optional). Simple network management protocol (SNMP) management stations (optional).
Obtain or assign intelligent port network addresses (optional).			SAN router intelligent ports require network addressing to support iFCP connectivity.
Obtain gateway addresses for router or other gateway devices on company LAN.			To configure management server and products (if installing on a LAN with non-McDATA devices).
Assign host names.			Management server and products (optional).
Add host name to domain name service (DNS) database.			Management server and products.
Determine what level of SAN management application user rights are to be used for up to 16 users.			

Table 6-3 Operational Setup Tasks (*continued*)

Activity	Task Owner	Due Date	Comments
Determine if inband management of the director or switch is to be used, and if so, the type (FICON or open-systems).			Management server and Fibre-Channel-attached server peripheral (optional).
Determine if the call-home feature is to be used.			
Determine call-home telephone numbers to be used.			
Determine if the e-mail notification feature is to be used.			Obtain e-mail addresses for event notification and identify e-mail server.
Determine SNMP access to directors and switches.			Obtain SNMP trap recipient IP addresses. Determine SNMP information required (generic and product-specific). Determine if write permission is required for modifying SNMP variables.
Determine if a multiswitch fabric is to be implemented.			Define the fabric topology (mesh, core-to-edge, or fabric (SAN) island).
Determine if SAN routing is to be implemented.			Define the distance extension operational mode and transport technology.
Determine if the zone management feature is to be used.			
Determine SAN routing zone policy (optional).			To support SAN routing, determine how zoning information is synchronized between a SAN router and attached fabrics.
Introduce staff to SAN management, SANvergence Manager, and Element Manager applications.			
Introduce staff to remote session parameters.			
Introduce staff to product recovery concepts and messages.			
Assign port names.			

Table 6-3 Operational Setup Tasks (*continued*)

Activity	Task Owner	Due Date	Comments
Configure extended distance ports.			If SAN routing is supported, configure extended distance ports in accordance with IP WAN requirements.
Enable and configure optional feature keys.			
Configure link incident alerts.			
Configure Ethernet events.			

A

Product Specifications

This appendix lists specifications for McDATA directors, fabric switches, storage area network (SAN) Routers, and the FC-512 Fabricenter equipment cabinet.

Director, Fabric Switch, and SAN Router Specifications

This section lists specifications (dimensions, weight, power requirements, heat dissipation requirements, cooling airflow clearances, acoustical noise generated, physical tolerances, storage and shipping environment requirements, and operating environment requirements) for directors, fabric switches, and SAN routers.

Dimensions

McDATA products have the following physical dimensions:

Intrepid 6064 Director:

Height: 39.7 centimeters (15.6 inches) or 9 rack units.

Width: 44.5 centimeters (17.5 inches).

Depth: 54.6 centimeters. (21.5 inches).

Weight: 53.1 kilograms (117.0 pounds).

Intrepid 6140 Director:

Height: 52.9 centimeters (20.8 inches) or 12 rack units.

Width: 44.5 centimeters (17.5 inches).

Depth: 61.3 centimeters. (24.1 inches).

Weight: 75.9 kilograms (167.0 pounds).

Intrepid 10000 Director:

Height: 62.2 centimeters (24.5 inches) or 14 rack units.

Width: 44.3 centimeters (17.5 inches).

Depth: 86.4 centimeters. (34.0 inches).

Weight: 152.0 kilograms (335.0 pounds).

Sphereon 3232 Fabric Switch:

Height: 6.5 centimeters (2.6 inches) or 1.5 rack units.

Width: 44.5 centimeters (17.5 inches).

Depth: 64.1 centimeters (25.2 inches).

Weight: 16.8 kilograms (37.0 pounds).

Sphereon 4300 Fabric Switch:

Height: 4.1 centimeters (1.6 inches) or 1 rack unit.

Width: 43.7 centimeters (17.2 inches).

Depth: 47.3 centimeters (18.6 inches).

Weight: 6.8 kilograms (15.0 pounds).

Sphereon 4400 Fabric Switch:

Height: 4.1 centimeters (1.6 inches) or 1 rack unit

Width: 19.9 centimeters (7.8 inches)

Depth: 33.3 centimeters (13.1 inches), plus 6.1 centimeters (2.4 inches) for external power supplies

Weight: 4.0 kilograms (8.8 pounds)

Sphereon 4500 Fabric Switch:

Height: 4.1 centimeters (1.6 inches) or 1 rack unit.

Width: 43.7 centimeters (17.2 inches).

Depth: 47.3 centimeters (18.6 inches).

Weight: 8.6 kilograms (19.0 pounds).

Sphereon 4700 Fabric Switch:

Height: 4.1 centimeters (1.6 inches) or 1 rack unit

Width: 43.7 centimeters (17.2 inches)

Depth: 39.4 centimeters (15.5 inches)

Weight: 6.8 kilograms (15.0 pounds)

Eclipse 1620 SAN Router:

Height: 4.1 centimeters (1.6 inches) or 1 rack unit.

Width: 43.7 centimeters (17.2 inches).

Depth: 45.7 centimeters (18.0 inches).

Weight: 5.9 kilograms (13.0 pounds).

Eclipse 2640 SAN Router:

Height: 4.1 centimeters (1.6 inches) or 1 rack unit.

Width: 43.7 centimeters (17.2 inches).

Depth: 68.6.7 centimeters (27.0 inches).

Weight: 10.9 kilograms (24.0 pounds).

Power Requirements

McDATA products have the following nominal power requirements:

Intrepid 6064 Director:

Input voltage: 100 to 240 VAC.

Input current: 2.0 amps at 208 VAC.

Input frequency: 47 to 63 Hz.

Intrepid 6140 Director:

Input voltage: 200 to 240 VAC.

Input current: 4.2 amps at 208 VAC.

Input frequency: 47 to 63 Hz.

Intrepid 10000 Director:

Input voltage: 180 to 270 VAC.

Input current: 12.0 amps at 208 VAC.

Input frequency: 47 to 63 Hz.

NOTE: The Intrepid 10000 Director must be connected directly to facility power outlets. The director draws a current of 32 amperes at power-on, and must not be connected to power strips in the FC-512 Fabricenter equipment cabinet.

Sphereon 3232 Fabric Switch:

Input voltage: 100 to 240 VAC.

Input current: 1.3 amps at 208 VAC.

Input frequency: 47 to 63 Hz.

Sphereon 4300 Fabric Switch:

Input voltage: 100 to 240 VAC.

Input current: 0.4 amps at 208 VAC.

Input frequency: 47 to 63 Hz.

Sphereon 4400 Fabric Switch:

Input voltage: 90 to 264 VAC

Input current: 0.18 amps at 208 VAC

Input frequency: 47 to 63 Hz

Sphereon 4500 Fabric Switch:

Input voltage: 100 to 240 VAC.

Input current: 0.5 amps at 208 VAC.

Input frequency: 47 to 63 Hz.

Sphereon 4700 Fabric Switch:

Input voltage: 90 to 264 VAC

Input current: 0.35 amps at 208 VAC

Input frequency: 47 to 63 Hz

Eclipse 1620 SAN Router:

Input voltage: 100 to 240 VAC.

Input current: 0.35 amps at 208 VAC.

Input frequency: 47 to 63 Hz.

Eclipse 2640 SAN Router:

Input voltage: 100 to 240 VAC.

Input current: 0.95 amps at 208 VAC.

Input frequency: 47 to 63 Hz.

Heat Dissipation

McDATA products have the following maximum heat dissipation characteristics:

Intrepid 6064 Director: 490 watts (1,672 BTU / hr).

Intrepid 6140 Director: 841 watts (2,873 BTU / hr).

Intrepid 10000 Director: 2,496 watts (8,517 BTU / hr).

Sphereon 3232 Fabric Switch: 245 watts (836 BTU / hr).

Sphereon 4300 Fabric Switch: 37 watts (127 BTU / hr).

Sphereon 4400 Fabric Switch: 36 watts (123 BTU / hr).

Sphereon 4500 Fabric Switch: 49 watts (167 BTU / hr).

Sphereon 4700 Fabric Switch: 71 watts (242 BTU / hr).

Eclipse 1620 SAN Router: 73 watts (249 BTU / hr).

Eclipse 2640 SAN Router: 198 watts (676 BTU / hr).

Clearances

McDATA products have the following cooling airflow clearances. In addition, the Intrepid 10000 Director may require removal from an equipment cabinet (left-side service clearance required) for FRU removal and replacement.

Intrepid 6064 and 10000 Directors:

Right and left side: 5.1 centimeters (2.0 inches).

Front and rear: 7.6 centimeters (3.0 inches).

Top and bottom: No clearance required.

NOTE: If the Intrepid 10000 Director is installed in a non-McDATA equipment cabinet with a door that does not provide direct airflow over the full height of the unit, 17.8 centimeters (7.0 inches) of front clearance is required.

Intrepid 6140 Director:

Right and left side: 2.5 centimeters (1.0 inches).

Front and rear: 7.6 centimeters (3.0 inches).

Top and bottom: No clearance required.

Sphereon 3232 Fabric Switch:

Right and left side: No clearance required.

Front and rear: 7.6 centimeters (3.0 inches).

Top and bottom: No clearance required.

Sphereon 4000-Series Switches:

Right and left side: 1.3 centimeters (0.5 inches).

Front and rear: 7.6 centimeters (3.0 inches).

Top and bottom: No clearance required.

Eclipse 1620 and 2640 SAN Routers:

Right and left side: No clearance required.

Front and rear: 7.6 centimeters (3.0 inches).

Top and bottom: No clearance required.

Acoustical Noise and Physical Tolerances

This section lists acoustical noise generated, shock and vibration tolerances, and inclination tolerances for McDATA directors, fabric switches, and SAN routers.

Acoustical noise generated:

Intrepid 6064 Director: 55 dB "A" scale.

Sphereon 4300 and 4500 Fabric Switches: 64 dB "A" scale.

Intrepid 10000 Director: 75 dB "A" scale.

All other products: 70 dB "A" scale.

Shock and vibration tolerance:

All products: 60 Gs for 10 milliseconds without nonrecoverable errors.

Inclination tolerance:

All products: 10^0 maximum.

Storage and Shipping Environment

This section specifies environmental requirements for storing and shipping McDATA products. Protective packaging must be provided for all domestic and international shipping methods.

Shipping temperature:

-40^0 F to 140^0 F (-40^0 C to 60^0 C).

Storage temperature:

34^0 F to 140^0 F (1^0 C to 60^0 C).

Shipping relative humidity:

5% to 100%.

Storage relative humidity:

5% to 80%.

Maximum wet-bulb temperature:

84^0 F (29^0 C).

Altitude:

40,000 feet (12,192 meters).

Operating Environment

This section specifies environmental requirements for operating McDATA products.

Temperature:

40^0 F to 104^0 F (4^0 C to 40^0 C).

Relative humidity:

8% to 80%.

Maximum wet-bulb temperature:

81^0 F (27^0 C).

Altitude:

10,000 feet (3,048 meters).

FC-512 Fabricenter Cabinet Specifications

This section lists specifications (dimensions, weight, power requirements, cooling airflow clearances, and service clearances) for the FC-512 Fabricenter equipment cabinet. An illustration of the cabinet footprint is also provided (Figure A-1).

Dimensions

The Fabricenter cabinet has the following physical dimensions:

Height: 186.1 centimeters (73.2 inches).

A total of 39 rack units (39 U) are available internal to the cabinet for product installation.

Width: 61.0 centimeters (24.0 inches).

Depth: 105.4 centimeters (41.5 inches).

Weight (no installed products): 153.2 kilograms (337.0 pounds).

Weight (shipping container): 50.9 kilograms (112.0 pounds).

Power Requirements

The Fabricenter cabinet has the following power requirements:

Input voltage: 200 to 240 VAC.

Input current: 30.0 amps at 208 VAC.

Input frequency: 47 to 63 Hz.

Clearances

The Fabricenter cabinet has the following cooling airflow and service clearances.

Cooling airflow clearances:

Right and left side: No clearance required.

Front and rear: 15.2 centimeters (6.0 inches).

Service clearances:

Right and left side: No clearance required.

Front and rear: 91.4 centimeters (36.0 inches).

Cabinet Footprint

Figure A-1 illustrates the Fabricenter cabinet footprint. The bottom denotes the front of the cabinet. The illustration includes:

1. Cabinet-leveling screws (four total).
2. Cooling airflow cutouts (eight total).
3. Fibre Channel cable cutouts (two total).
4. Caster wheel attachments (four total).
5. Power cable cutout (one).

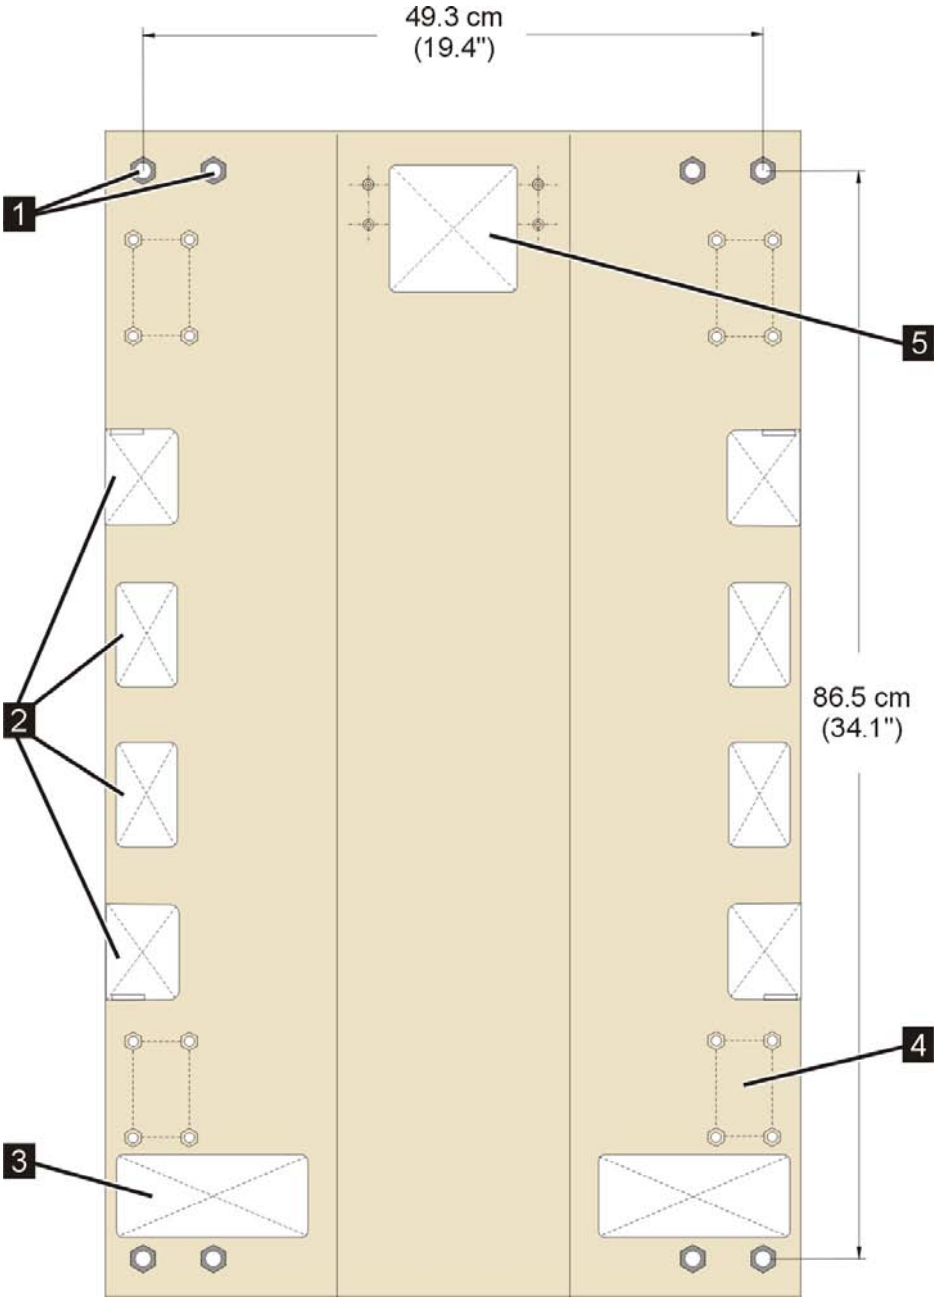

Figure A-1 Fabricenter Cabinet Footprint

B

Firmware Summary

This appendix summarizes differences and similarities between the Enterprise Operating System, classic (E/OSc) for Intrepid 6000-series directors and Sphereon-series fabric switches; Enterprise Operating System, nScale (E/OSn) for the Intrepid 10000 Director; and Enterprise Operating System, internetworking (E/OSi) for Eclipse-series SAN routers. The appendix includes tables that list:

- System-related differences.
- Fibre Channel protocol-related differences.
- Management-related differences.

System-Related Differences

Table B-1 summarizes system-related differences between the E/OSc 8.0, E/OSn 6.0, and E/OSi 4.6 firmware versions.

Table B-1 E/OSc versus E/OSn and E/OSi - System-Related Differences

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
Operating system	E/OSc with VxWorks® and McDATA-proprietary kernels.	E/OSn with MontaVista™ Linux (CTP cards) and ThreadX® line module (LIM) kernels.	E/OSi with VxWorks® with McDATA-proprietary modifications.
Initial program load (IPL) requirements	IPL required to activate a product feature enablement (PFE) key or reset software.	IPL functionality not available. PFE key activation does not require code restart.	IPL functionality not available. PFE key activation not applicable.

Table B-1 E/OSc versus E/OSn and E/OSi - System-Related Differences (Continued)

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
Nondisruptive hot code activation (HotCAT)	Upgrade nondisruptive to Fibre Channel traffic for single and dual CTP card directors and for fabric switches. Upgrade does not require director CTP card switchover. E/OSc allows upgrade or downgrade by greater than one functional release.	Director must be set offline to upgrade a single CTP card. Upgrade nondisruptive to Fibre Channel traffic for dual CTP card but requires card switchover. New firmware image provided for all CTP cards and eight LIMs. E/OSn allows upgrade or downgrade by only one functional release.	Nondisruptive hot code activation not applicable.
Initial microcode load (IML) functionality	IML functionality is supported.	IML functionality not available.	IML functionality not available.
Non-redundant CTP card upgrade (directors only)	Concurrent CTP card upgrade supported.	CTP card upgrade disruptive to Fibre Channel traffic and requires director to be set offline.	CTP card upgrade not applicable.
Switching logic card function (directors)	One active serial crossbar assembly (SBAR) and one standby SBAR supported.	Four load-sharing active switching module (SWM) cards supported, eliminating the need for standby modules.	Switching logic card function not applicable.
Optical paddles (directors)	Optical paddles not applicable. Port cards provide equivalent functionality.	Concurrently replaceable optical paddles are supported, but Fibre Channel traffic is disrupted to all ports on the paddle.	Optical paddles not applicable.
Power supplies	Two power modules (or power supplies) with two power cords. The Sphereon 4300 Fabric Switch has one power supply and one power cord.	Four power modules with two power cords.	Two power supplies with two power cords. Power cords connect to the front panel of the Eclipse 1620 SAN Router and to the back of the Eclipse 2640 SAN Router.
Power-on diagnostic step (P-step) codes	Numeric format.	Text-based format.	Text-based format.
Power-on-hour (POH) updates	POH vital product data (VPD) updated to field-replaceable unit (FRU) read-only memory (ROM) every hour.	POH VPD updated to CTP card memory every hour then downloaded to FRU ROM upon FRU failure.	Power-on-hour updates not applicable.
10.2000 Gbps port transmission speed support	10.2000 Gbps port operation not available.	10.2000 Gbps port operation supported.	10.2000 Gbps port operation not available.

Table B-1 E/OSc versus E/OSn and E/OSi - System-Related Differences (Continued)

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
Misaligned word generation	Misaligned words not generated.	Misaligned word is generated when the port state machine (PSM) transitions from inactive to active state or from active to inactive state. This anomaly has no effect on Fibre Channel device behavior.	Misaligned words not generated.
Flexible partition (FlexPar) feature support	Flexpar feature not supported.	Up to four FlexPars supported for each Intrepid 10000 Director chassis.	Flexpar feature not supported.
Maintenance port access	For directors, one 9-pin, RS-232 serial port shared between two CTP cards. Port-attached device communicates only with the active card. For fabric switches, one 9-pin, RS-232 serial port per switch.	Each CTP card has an independent 9-pin, RS-232 serial port that provides access to engineering-level interfaces. Maintenance mode and command line interface (CLI) available through the active CTP card port.	One 9-pin, RS-232 serial port per SAN router.

Fibre Channel Protocol-Related Differences

Table B-2 summarizes Fibre Channel protocol-related differences between the E/OSc 8.0, E/OSn 6.0, and E/OSi 4.6 firmware versions.

Table B-2 E/OSc versus E/OSn and E/OSi - Fibre Channel Protocol-Related Differences

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
Node port (N_Port) ID assignment	N_Port ID assigned based on port number plus offset of four. Device attached to port ***nn*** obtains N_Port ID of ***dd(nn+4)xx***, where ***dd*** is the switch domain ID and ***xx*** is hexadecimal 13.	N_Port ID assigned based on port number. Device attached to Port ***nn*** obtains N_Port ID of ***ddnnxx***, where ***dd*** is the switch domain ID and ***xx*** is hexadecimal 13.	N_Port ID assigned based on port number. Device attached to Port ***nn*** obtains N_Port ID of ***ddnnxx***, where ***dd*** is the switch domain ID (always 01) and ***xx*** is hexadecimal 13.
Exchange switch support (ESS) sequence transmission	ESS not transmitted until fabric shortest path first (FSPF) algorithm obtains best hop to target domain.	ESS transmitted at the moment adjacency to the target domain is detected (link at FSPF FULL state). Therefore, ESS always transmitted sooner for an Intrepid 10000 Director.	ESS transmitted only in response to an ESS received from another fabric element.

Table B-2 E/OSc versus E/OSn and E/OSi - Fibre Channel Protocol-Related Differences

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
Expansion port (E_Port) staging	E_Port staging not supported.	The Intrepid 10000 Director supports E_Port staging. The director allows only one E_Port connection to a neighbor switch to take part in a fabric build process. Upon process completion, subsequent E_Ports to neighbor switches are brought up (staged). In addition, staged E_Ports are brought up in groups to minimize Class F traffic and avoid overloading.	E_Port staging not supported.
ESS payload processing and domain port number zoning	E/OSc Version 6.0 (or earlier) handshakes the ESS sequence but does not apply content of ESS payload to calculate port numbering (PN) to logical port address (PA) mapping. Legacy directors and switches are assumed to use an offset of four for PN to PA mapping.	The Intrepid 10000 Director sends and accepts payloads specifying minimum and maximum logical port addresses and processes payload content to calculate PN to PA mapping. A correct mapping scheme interprets domain port number zoning enforcement across the fabric and allows the director to participate in a legacy fabric with domain port number zoning.	E/OSi handshakes the ESS sequence. However, domain port number zoning not supported.
Registered state change notification (RSCN) coalescing	Only fabric login (FLOGI) and Name Server events are coalesced.	The Intrepid 10000 Director coalesces all events that trigger and transmit RSCNs to N_Ports. This minimizes the number of RSCNs generated.	SAN routers coalesce all RSCN events.
Generating fabric format RSCNs after CTP failover (directors)	Does not replicate RSCN event triggers, therefore transmits a fabric format RSCN to all attached devices after CTP card failover.	Replicates events that trigger RSCNs to the backup CTP, therefore does not generate a fabric format RSCN after CTP failover.	RSCNs after CTP failover not supported.
Hop count restriction	Hop count of up to three is supported. Devices more than three hops away cannot be reached.	No hop count restriction is applied.	No hop count restriction is applied.
Exchange switch capabilities (ESC) support	SW_ILS ESC sequence not supported.	The Intrepid 10000 Director processes SW_ILS ESC sequences to identify neighboring director ports.	SW_ILS ESC sequence not supported.

Table B-2 E/OSc versus E/OSn and E/OSi - Fibre Channel Protocol-Related Differences

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
Reroute delay behavior	With reroute delay enabled, the destination route point is cleared and a delay equal to the error detect time-out value (ED_TOV) is applied before a new route is programmed. During the delay, Class 3 frames are dropped.	With reroute delay enabled, the Intrepid 10000 Director first pauses incoming port traffic and applies a delay to allow frames internal to the director to be transmitted. This delay during route reprogramming prevents frames being sent out of order.	Reroute delay not supported.
Switch internal link services (SW_ILS) during fabric build	SW_ILS sequences transmitted on up to eight interswitch links (ISLs) per neighbor switch.	SW_ILS sequences transmitted only on the primary ISL to a neighbor switch. However, if the domain identifier assigned (DIA) link service is enabled, the Intrepid 10000 Director floods all ISLs. This reduces Class F traffic in a large fabric.	SW_ILS sequences not supported.

Management-Related Differences

Table B-3 summarizes management-related differences between the E/OSc 8.0, E/OSn 6.0, and E/OSi 4.6 firmware versions.

Table B-3 E/OSc versus E/OSn and E/OSi - Management-Related Differences

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
Command line interface (CLI) scope and syntax	CLI supports a configuration-related subset of product features.	Exhaustive CLI command set supported, including all configurable features of the Intrepid 10000 Director. Different in syntax and semantics from E/OSc and E/OSi CLIs.	Exhaustive CLI command set supported, including all configurable features of SAN routers. Different in syntax and semantics from E/OSc and E/OSn CLIs.
EFCM Basic Edition interface	HTML-based EFCM Basic Edition interface supported.	EFCM Basic Edition interface not supported.	EFCM Basic Edition interface not supported. Embedded Element Manager application is a Java applet, not HTML-based.
Proprietary MIB support	Proprietary MIB interface supports a configuration-related subset of product features.	Proprietary MIB interface exhaustive and supports all configurable features of the Intrepid 10000 Director.	Proprietary MIB interface exhaustive and supports all configurable features of SAN routers.

Table B-3 E/OSc versus E/OSn and E/OSi - Management-Related Differences (Continued)

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
General SNMP support	SNMP interface supports Version 1, Version 2c, and a configuration-related subset of product features.	SNMP interface supports Version 1, Version 2c, and Version 3. Read-only use recommended with Versions 1 and 2c. Read-write use recommended with secure Version 3.	SNMP interface supports Version 1 and a configuration-related subset of product features.
Fibre Alliance management information base (MIB) support	Supports Fibre Channel Management Framework Integration MIB (FC-MGMT-MIB) Versions 3.0 and 3.1. MIB version user defined.	Supports FC-MGMT-MIB Version 4.0. MIB version not user defined.	Supports FC-MGMT-MIB Version 3.0. MIB version not user defined.
SNMP traps	Subset of event notifications sent as SNMP traps. All transmissions use proprietary not most recently used (NMRU) protocol.	All significant events asynchronously distributed as SNMP traps.	Subset of event notifications asynchronously distributed as SNMP traps.
Throughput threshold alert (TTA) support	Port TTAs supported through the EFCM and CLI interfaces. Switch performance threshold alerts (SPTAs) not supported.	Port TTAs supported through the Intrepid 10000 Director CLI and EFCM interfaces. Paddle-pair SPTAs only supported through the director CLI interface.	Port TTAs and SPTAs not supported.
Counter threshold alert (CTA) support	Port counter CTAs supported through the CLI interface only.	Port counter CTAs not supported.	Port counter CTAs not supported.
Product feature enablement (PFE) key support	PFE keys supported: Element Manager, OpenTrunking, SANtegrity binding and authentication, OSMS, FMS, FlexPorts, FICON CUP zoning, full volatility, full fabric, and NPIV.	PFE keys supported: Element Manager, SANtegrity binding, FMS, remote fabric, and FlexPars (last two are Intrepid 10000 Director features). Full volatility is supported but does not require a PFE key.	PFE keys are not supported through the SANvergence Manager or Element Manager applications.
Port address **FE** and **FF** prohibit dynamic connectivity mask (PDCM) array entries	Port addresses **FE** and **FF** are implemented internally and only when FMS is PFE-key enabled.	Port addresses **FE** and **FF** do not exist in the PDCM array and cannot be used. FICON devices cannot attach to physical ports associated with these addresses. When FMS is enabled, the ports become spare (transmit only an OLS) but can be port swapped.	FICON management style (with associated PDCM array) not supported.

Table B-3 E/OSc versus E/OSn and E/OSi - Management-Related Differences (Continued)

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
FICON Management Server (FMS) PFE key install and enable behavior	FMS is auto-enabled upon PFE installation. Management style automatically changes to FICON.	When PFE is installed, the user must explicitly enable FMS. Management style does not automatically change to FICON.	FMS PFE key not supported.
Saving a user-defined PDCM array through EFCM	Saving a user-defined PDCM array through EFCM supported.	Cannot save a user-defined PDCM array through EFCM.	FICON management style (with associated PDCM array) not supported.
Saving a PDCM array after reset (single CTP card)	No limitation. When a single-CTP card director is reset, the PDCM array is saved.	When *Active = Saved* is not set and a single-CTP card Intrepid 10000 Director is reset, the PDCM array is not saved. This does not apply to a redundant CTP card configuration.	FICON management style (with associated PDCM array) not supported.
Full volatility support	Full volatility supported through PFE key.	Full volatility always supported because the Intrepid 10000 Director does not persistently store data frame contents on any FRU.	Full volatility not supported.
Port configuration behavior as a function of port module type (directors)	Configuration set based on last inserted port card type (FPM, UPM, or XPM). If a new port card is different in type from the last inserted card, then all port-level parameters for the new card are reset to default. Parameters include port type, transmission speed, BB_Credit, block or unblock, preferred path, allow or prohibit, and port swap settings.	Configuration set independent of port card type, except transmission speed and BB_Credit are configured and stored independently for 2.1250 and 10.2000 Gbps ports (user specifies port number and type). Port-level parameters are not reset to default when a new port card is inserted.	Port configuration behavior as a function of port module type not applicable.
Display of Fibre Channel link speed	Fibre Channel link speed always displayed, regardless of switch state.	Fibre Channel link speed not displayed when Intrepid 10000 Director is disabled.	Configured link speed displayed at Element Manager application main window and actual link speed displayed at *Port Configuration* dialog box, regardless of switch state.
Port speed configuration	Ports are automatically set offline or offline and not required to be blocked to set transmission speed.	Ports required to be blocked (offline) to set transmission speed.	Ports are not required to be set offline to set transmission speed.

Table B-3 E/OSc versus E/OSn and E/OSi - Management-Related Differences (Continued)

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
Port group online diagnostics (directors)	Diagnostic errors for any single port in a group (port card) cause a group failure indication.	Diagnostic errors for any single port in a group (LIM) are indicated individually and do not cause a group failure indication.	Port grouping (port cards or LIMs) not supported.
Online state behavior	Product online or offline state not persistent. An IML or IPL sets the product online.	Intrepid 10000 Director online or offline state is persistent.	Online or offline state set at the port level and is persistent only if setting is saved to FLASH memory.
Port reset while blocked	Port is reset independent of port blocking state.	Port is reset only if port is not blocked.	Port reset supported for Eclipse 2640 intelligent ports (**13** to **16**), independent of port blocking state.
Requirements to set product offline	Product can be online to set domain RSCN state and zone change RSCN control state.	Intrepid 10000 Director must be offline to set domain RSCN state and zone change RSCN control state.	Online or offline status to set domain RSCN state or zone change RSCN control state not applicable.
Ability to disable management interfaces	Can selectively disable SNMP and CLI management interfaces.	Cannot selectively disable SNMP and CLI management interfaces.	Cannot selectively disable SNMP and CLI management interfaces.
Port technology information	Port technology (optical transceiver) information available when the product is online or offline.	Port technology (optical transceiver) information available only when the Intrepid 10000 Director is online.	Port technology (optical transceiver) information not available.
Port type and speed reporting (port offline or blocked)	Offline or blocked ports report the configured port type and speed, not the actual port type and speed.	Offline or blocked ports report **G_Port** and **Not Established** for the actual port type and speed. When the Intrepid 10000 Director is set offline, port state changes are not reflected at the user interface.	Online and offline ports report configured port type and speed at Element Manager application main window and actual port type and speed at *Port Configuration* dialog box. No difference between configured and actual port type.
Port type displayed in the *Port List View*.	If no established connection or cable attached to port transceiver, *Type* field entry defaults to configured port type.	If no established connection or cable attached to port transceiver, *Type* field entry defaults to **Unknown**.	*Port List View* not supported by the SANvergence Manager application. However, port type is always displayed as configured, even when port is offline.

Table B-3 E/OSc versus E/OSn and E/OSi - Management-Related Differences (Continued)

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
Special port states and reason codes	Port state **Inactive** with special reason codes **Reserved** and **InvalidOTPConfig** not supported.	Logical port addresses **FE** and **FF** forced offline and placed in special state **Inactive** with reason code **Reserved** if FICON CUP is enabled. FICON states are **Internal Port** (address **FE**) and **Unimplemented** (address **FF**). If the top or bottom optical paddle pairs in a line module are both 10.2000 Gbps paddles, then ports in one paddle are set offline and placed in the **Inactive** state with special reason code **InvalidOTPConfig**. If the director powers up with two 10.2000 Gbps paddles in a pair, the upper paddle is set offline. If a second 10.2000 Gbps paddle is added to an operating director, the newly-inserted paddle is set offline.	At the CLI, port state **Testing** indicates a Fibre Channel port is enabled without a device attached. At the CLI and Element Manager application, port state **Needs Reboot** indicates SAN Router requires a reboot because a Fibre Channel port is configured as a GbE port (or vice versa).
Inactive zone set support	Only the active zone set is stored.	Up to 32 inactive zone sets saved and supported through the CLI interface.	Only the active zone set is automatically stored. Selected inactive zone sets saved and supported through the SANvergence Manager application.
Maintenance mode command set	Maintenance mode commands supported through the product serial maintenance port (protected access), but command set differs from the Intrepid 10000 Director.	Maintenance mode commands supported through the product serial maintenance port (protected access). Maintenance mode entered as a special user through the CLI interface. All CLI commands available in maintenance mode.	Maintenance mode commands supported through the product serial maintenance port (protected access) or Telnet connection.
Diagnostics (port granularity)	Diagnostics supported for one port or all ports on FPM, UPM, or XPM cards.	Diagnostics supported for one port or all ports on an optical paddle pair. User provides a data pattern and duration through the CLI interface.	Diagnostics supported for one port or all ports when the SAN router is offline.

Table B-3 E/OSc versus E/OSn and E/OSi - Management-Related Differences (Continued)

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
Diagnostics (port failure state)	Port set to **Failed** state upon failing a user- invoked port diagnostic test.	If external port diagnostics are performed without a loopback plug the diagnostic test fails, but port is not set to **Failed**.	If external port diagnostics or management port diagnostics are performed without a loopback plug, the diagnostic test fails, but port is not set to **Failed**.
Buffer-to-buffer credit (BB_Credit) support	BB_Credit subject to a per-port maximum value and a product-wide (pool) maximum value. Refer to *Distance Extension Through BB_Credit* for information.	BB_Credit subject to a per-port maximum value and a paddle-pair wide (pool) maximum value. BB_Credit also subject to remote fabric PFE being enabled or disabled. Refer to *Distance Extension Through BB_Credit* for information.	Fibre Channel ports set to a BB_Credit value of 16. Intelligent ports support IP network connectivity and not assigned BB_Credits. Ports provide 96 MB of TCP packet buffering per transmission direction.
Audit log entries related to flexible partitioning	Flexpar feature not supported.	When a user downloads firmware, performs a CTP card switchover, configures partitioning, or clears the system error light (or performs an operation with system-wide impact), the corresponding *Audit* log entry is recorded only for the **ADMIN** partition (**0**). Log entries are processed differently for the **ADMIN** partition (**0**) versus other partitions (**1** through **3**).	Flexpar feature not supported.
Other log entries related to flexible partitioning	Flexpar feature not supported.	Log entries (non *Audit* log) are processed differently for the **ADMIN** partition (**0**) versus other partitions (**1** through **3**).	Flexpar feature not supported.
Audit log entries (EFCM)	Log contains EFCM user name and IP address of EFCM client.	Log contains only the IP address of the EFCM server. Log does not contain EFCM user name and IP address of EFCM client.	Remote workstations (clients) not supported through the SANvergence Manager application. The application supports an APP log that is functionally equivalent to the EFCM *Audit* log. In addition, the Element Manager application has an *Audit* log independent of SANvergence Manager.

Table B-3 E/OSc versus E/OSn and E/OSi - Management-Related Differences (Continued)

Feature	E/OSc 8.0	E/OSn 6.0	E/OSi 4.6
FICON management style support	User-selectable option on a per product basis. Setting is stored on the product.	User-selectable option on a per user (different for each login ID) or per product basis. Setting is stored on the management server, not the product. Setting is backed up by the server backup process.	FICON management style not supported through the SANvergence Manager application.
LED status indicators	For director CTP and SBAR card LEDs: Green LED indicates active or standby. Amber LED indicates failed or beaconing. For director FPM, UPM, or XPM card LEDs: Amber LED indicates failed or beaconing. For switches: Green **PWR** LED indicates operational. Amber **ERR** LED indicates failed or beaconing.	For CTP, SWM, or LIM card **State** LED (bicolor LED): Green indicates online. Amber indicates failed, degraded, or beaconing. For CTP card **Role** LED (bicolor LED): Green indicates active or standby. Amber indicates failed.	Green **SYS** LED indicates SAN router operational. For Eclipse 1620 SAN Router, green **PS1** and **PS2** LEDs indicate power supplies operational.
Access to simple name server (SNS) database	User can only view local SNS database through the EFCM application and CLI interface.	User can view local SNS database through the EFCM application and CLI interface. User can view global SNS databases only through the CLI interface.	User can view local and global SNS databases through the SANvergence Manager application. User can view only local SNS database through the Element Manager application. CLI interface only shows SNS database for direct-attached devices.
Date and time synchronization	Periodic and user-initiated *Sync Now* date and time synchronization between EFCM application and product supported.	Periodic date and time synchronization between EFCM application and product not supported. Application supports user-initiated *Sync Now* function for the Intrepid 10000 Director.	Periodic date and time synchronization between SANvergence Manager application and product not supported.
Hard zoning restriction (fabric size and port type)	No zoning restriction based on fabric size and port type.	When a paddle pair exceeds seven E_Port connections and the fabric node count exceeds 800 devices, hard zoning is disabled for that paddle pair.	No zoning restriction based on fabric size and port type.

Index

A

access control list
- inband 5-18
- out-of-band 5-18
- role-based access 4-7

any-to-any connectivity 1-28
application I/O profiles 3-31
arbitrated loop switch
- connectivity features 1-28
- default network address 6-9
- security features 1-31
- serviceability features 1-32

arbitrated loop topology
- description 3-2
- fabric-attached planning considerations 3-11
- FL_Port connectivity 3-10
- operating characteristics 3-2
- private device connectivity 3-7
- private loop 3-9
- private loop planning considerations 3-10
- public device connectivity 3-6
- public loop 3-8
- shared mode operation 3-3
- switched mode operation 3-4

architecture
- E/OSc description 2-11
- E/OSi description 2-12
- E/OSn description 2-11
- EON 1-9, 1-10
- firmware differences B-1
- firmware similarities B-1
- nScale 1-12

asynchronous remote data replication
- description 4-38
- IP transport 4-48
- long-distance requirement 4-38
- SONET/SDH transport 4-47

B

backup features
- CD-RW drive 2-13
- director and fabric switch NV-RAM configuration 2-13
- SAN management data directory 2-14
- SAN router NV-RAM configuration 2-14

bandwidth
- dark fiber transport 4-46
- dedicated (SAN routing) 4-55
- director ports 1-7
- fabric switch ports 1-15
- IP transport 4-48
- IRL 4-25
- ISL 3-22
- rate limiting (SAN routing) 4-51
- SAN router ports 1-23
- SONET/SDH transport 4-47
- WDM transport 4-47

BB_Credit
- configure at Element Manager 1-30, 4-50
- description 1-30, 4-49
- extended distance support 1-29
- full fabric feature 5-38
- Intrepid 10000 Director 4-50
- remote fabric feature 5-39
- Sphereon 3232 Fabric Switch 1-29

Sphereon 4300 Fabric Switch 1-30
Sphereon 4400 Fabric Switch 1-30
Sphereon 4500 Fabric Switch 1-30
Sphereon 4700 Fabric Switch 1-30
best practices
cabling 3-20
configuration planning 6-1
connectivity 3-20
distance extension 4-55
FCP and FICON intermix 3-47
FICON cascading 3-54
multiswitch fabric topology 3-20
preventing ISL oversubscription 3-33
SAN routing 4-29
security 5-30
binding
fabric 5-19
persistent 5-29
SANtegrity 5-19
switch 5-19
business continuance
IP versus storage traffic 4-36
operational mode
asynchronous remote data replication 4-38
synchronous remote data replication 4-38
requirements
data priority 4-37
distance 4-37
recovery point objective 4-37
recovery time objective 4-37
transport technology
dark fiber 4-39
IP 4-45
SONET/SDH 4-42
WDM 4-40
transport technology comparison 4-48

C

cable routing
Ethernet requirements 5-11
fiber-optic requirements 5-8
planning considerations 6-3
cabling
50/125 multimode 5-7
62.5/125 multimode 5-7
9/125 singlemode 5-7
best practices 3-20
Ethernet cable routing 5-11
fiber-optic cable routing 5-8
planning considerations 6-3
port requirements 5-2
call-home support
description 1-33
telephone connection 6-12
CD-RW drive 2-13
checklist
operational setup tasks 6-36
planning and hardware installation tasks 6-35
class of service
Class 2 1-8
Class 3 1-8
Class F 1-8
clearances
directors A-5
fabric switches A-5
Fabricenter cabinet A-8
SAN routers A-5
command line interface 2-21
configuration diagram 6-12
connectivity
best practices 3-20
features 1-28
SAN router (logical) 4-12
SAN router (physical) 4-11
connector, LC duplex 5-7
consolidating
iSCSI servers 4-60
iSCSI storage 4-61
SAN islands
FlexPar technology 4-4
SAN routing 4-8
servers 3-12
tape devices 3-13
core-to-edge fabric
description 3-16
illustration, 2-by-14 3-17
suitability 3-16

Tier 1 connections 3-17
Tier 2 connections 3-18
Tier 3 connections 3-18
CT authentication 5-17

D

dark fiber distance extension
bandwidth 4-46
description 4-39
illustration 4-40
latency 4-46
recovery point objective 4-46
recovery time objective 4-46
data compression
algorithm selections 4-26
description 1-24
optimizing WAN use 4-55
set compression level 4-57
data transmission distance
cable type 5-4
multiswitch fabric requirements 3-21
transceiver type 5-4
default network addresses
director or fabric switch 6-9
Eclipse 1620 SAN Router 6-9
Eclipse 2640 SAN Router 6-9
management server 6-8
device
connectivity
Tier 1 3-17
Tier 2 3-18
Tier 3 3-18
fan-out ratio 3-34
locality 3-33
device window
description 2-19
illustration 2-19
dimensions
directors A-1
fabric switches A-1
Fabricenter cabinet A-8
SAN routers A-1
director
connectivity features 1-28
default network address 6-9
description 1-6
FlexPars 4-4
performance features 1-7
product overview 1-2
security features 1-31
serviceability features 1-32
specifications A-1
disaster recovery
IP versus storage traffic 4-36
operational mode
asynchronous remote
data replication 4-38
synchronous remote
data replication 4-38
requirements
data priority 4-37
distance 4-37
recovery point objective 4-37
recovery time objective 4-37
transport technology
dark fiber 4-39
IP 4-45
SONET/SDH 4-42
WDM 4-40
transport technology comparison 4-48
distance extension
assigning BB_Credits 4-50, 5-6
best practices 4-55
full fabric feature 5-38
operational mode
asynchronous remote
data replication 4-38
synchronous remote
data replication 4-38
port configuration 5-6
remote fabric feature 5-39
support 1-29
transport technology
dark fiber 4-39
IP 4-45
SONET/SDH 4-42
WDM 4-40
transport technology comparison 4-48
Domain_ID assignment
director 3-24
fabric switch 3-24

proxy Domain_ID 30 4-13, 4-18
proxy Domain_ID 31 4-13, 4-25
R_Port 4-15
SAN router 4-15

E

E/OSc
description 2-11
management-related properties B-5
operating system differences B-1
operating system similarities B-1
protocol-related properties B-3
system-related properties B-1
E/OSi
description 2-12
management-related properties B-5
operating system differences B-1
operating system similarities B-1
protocol-related properties B-3
system-related properties B-1
E/OSn
description 2-11
management-related properties B-5
operating system differences B-1
operating system similarities B-1
protocol-related properties B-3
system-related properties B-1
E_Port
DHCHAP authentication 5-17
full fabric feature 1-17
port fencing 1-36
segmentation 3-27
Eclipse 1620 SAN Router
default network address 6-9
description 1-24
FRUs 1-26
iFCP protocol 4-23
illustration 1-25
intelligent ports 1-25
Eclipse 2640 SAN Router
default network address 6-9
description 1-26
FRUs 1-27
iFCP protocol 4-23, 4-61
illustration 1-26
intelligent ports 1-27
mFCP protocol 4-20
EFCM application
description 2-15
GUI description 2-15
product overview 1-5
EFCM Basic Edition interface
description 2-20
plan console support 6-6
server connectivity 5-14
EFCM Lite application
description 2-3
unsupported features 2-4
Element Manager application
description (director and fabric switch) 2-16
description (SAN router) 2-19
feature key description 5-39
Hardware view 2-16
e-mail notification
description 1-33
support planning 6-11
Enterprise Fabric mode 5-19
environment
operating A-7
shipping A-6
storage A-6
EON architecture 1-9, 1-10
Ethernet cabling
access planning 6-7
management server 5-10
remote workstations 5-11
routing 5-11
Ethernet hub
description 2-10
illustration 2-10
extended distance
assigning BB_Credits 4-50, 5-6
best practices 4-55
full fabric feature 5-38
operational mode
asynchronous remote data replication 4-38
synchronous remote data replication 4-38

port configuration 5-6
remote fabric feature 5-39
support 1-29
transport technology
dark fiber 4-39
IP 4-45
SONET/SDH 4-42
WDM 4-40
transport technology comparison 4-48

F

fabric availability
nonresilient dual fabric 3-38
nonresilient single fabric 3-37
redundant fabrics 3-38
resilient dual fabric 3-38
resilient single fabric 3-38
fabric binding 5-19
fabric element
FCP and FICON intermix environment 3-41
limitations in a fabric 3-19
fabric performance
device fan-out ratio 3-34
fabric initialization 3-29
fabric scalability 3-39
I/O requirements 3-31
performance tuning 3-35
fabric switch
connectivity features 1-28
default network address 6-9
description 1-14
performance features 1-15
product overview 1-2
security features 1-31
serviceability features 1-32
specifications A-1
Fabricenter cabinet
footprint A-8
illustration 1-4
product overview 1-3
specifications A-7
fan-out ratio 3-34
FastWrite technology
description 1-24
IRL optimization 4-27
optimizing WAN use 4-56
FC-AL devices
connecting to a multiswitch fabric 3-11
server consolidation 3-12
tape device consolidation 3-13
FCP and FICON intermix
best practices 3-47
fabric element management 3-41
impacting features 3-46
management limitations 3-44
port numbering and logical port addressing 3-42
feature key
description 5-33
Element Manager application 5-39
Flexport Technology 5-36
FMS 5-35
format 5-34
full fabric 5-38
full volatility 5-38
N_Port ID virtualization 5-39
OpenTrunking 5-37
OSMS 5-35
remote fabric 5-39
SANtegrity Authentication 5-36
SANtegrity Binding 5-36
fiber-optic cabling
50/125 multimode 5-7
62.5/125 multimode 5-7
9/125 singlemode 5-7
overview 5-1
planning considerations 6-3
routing 5-8
Fibre Channel topology
arbitrated loop topology 3-2
core-to-edge fabric 3-16
mesh fabric 3-14
multiswitch fabric topology 3-2
SAN island 3-18
FICON cascading
best practices 3-54
definition 3-47
general description 3-52
high-integrity fabrics 3-53
minimum requirements 3-53

FICON management server
description 5-35
introduction 2-5
plan console support 6-6
firmware
application services 2-12
E/OSc description 2-11
E/OSi description 2-12
E/OSn description 2-11
fabric services 2-13
Fibre Channel protocol services 2-12
loop services 2-13
network services 2-12
operating system differences B-1
operating system services 2-12
operating system similarities B-1
port services 2-12, 2-13
system management services 2-12
FL_Port connectivity 3-10
FlexPar technology
description 4-4
director FlexPars 4-4
inter-FlexPar routing 4-28
Intrepid 10000 Director 1-12, 4-4
master FlexPar 4-5
role-based FlexPars 4-7
SAN island consolidation 4-4
zone FlexPars 4-6
Flexport Technology feature
description 5-36
Sphereon 3232 Fabric Switch 1-16
Sphereon 4300 Fabric Switch 1-17
Sphereon 4400 Fabric Switch 1-19
Sphereon 4500 Fabric Switch 1-20
Sphereon 4700 Fabric Switch 1-22
FMS feature
description 5-35
introduction 2-5
plan console support 6-6
footprint, Fabricenter cabinet A-8
frame delivery order 3-26
FRUs
Eclipse 1620 SAN Router 1-26
Eclipse 2640 SAN Router 1-27
Intrepid 10000 Director 1-13
Intrepid 6064 Director 1-9
Intrepid 6140 Director 1-10
Sphereon 3232 Fabric Switch 1-15
Sphereon 4300 Fabric Switch 1-17
Sphereon 4400 Fabric Switch 1-18
Sphereon 4500 Fabric Switch 1-20
Sphereon 4700 Fabric Switch 1-21
full fabric feature
description 5-38
Sphereon 4300 Fabric Switch 1-17
full volatility feature 5-38

G

gateway address
director or fabric switch 6-9
Eclipse 1620 SAN Router 6-9
Eclipse 2640 SAN Router 6-10
management server 6-8
graphical user interface
device window 2-19
EFCM application 2-15
EFCM Basic Edition 2-20
Hardware view 2-17
main window (SAN router) 2-18
SANavigator application 2-15
SANvergence Manager application 2-18

H

Hardware view
description 2-16
illustration 2-17
heat dissipation
directors A-5
fabric switches A-5
SAN routers A-5
high-availability
directors 1-7
fabric switches 1-15
planning considerations 5-6
SAN routers 1-24
high-integrity fabrics 3-53
hop count
limitations 3-20
minimizing (SAN routing) 4-56
path selection 3-25
HotCAT technology 1-9, 1-10, 1-12

I

I/O requirements
- application I/O profiles 3-31
- device locality 3-33
- ISL oversubscription 3-32

iFCP protocol
- build fabric events 4-23
- comparison to mFCP 4-28
- description 4-22

inband product management
- feature keys 5-35
- FMS feature 2-5
- OSMS feature 2-5
- plan console support 6-6

intelligent port
- Eclipse 1620 SAN Router 1-25
- Eclipse 2640 SAN Router 1-27
- implement rate limiting 4-51
- set port speed 4-52

inter-FlexPar routing 4-28

interoperability
- planning 6-4
- vendor limitations 3-20

Intrepid 10000 Director
- default network address 6-9
- description 1-12
- extended distance support 1-29
- FlexPar technology 1-12, 4-4
- FRUs 1-13
- illustration 1-12
- large fabric support 3-30
- LIMs 5-2

Intrepid 6064 Director
- default network address 6-9
- description 1-8
- FRUs 1-9
- illustration 1-9
- port cards 5-2

Intrepid 6140 Director
- default network address 6-9
- description 1-10
- FRUs 1-10
- illustration 1-11
- port cards 5-2

IP address
- director or fabric switch 6-9
- Eclipse 1620 SAN Router 6-9
- Eclipse 2640 SAN Router 6-9
- management server 6-8

IP distance extension
- bandwidth 4-48
- description 4-45
- illustration 4-46
- latency 4-48
- recovery point objective 4-48
- recovery time objective 4-48

IRL
- description 4-9
- optimization
  - data compression 4-26
  - FastWrite technology 4-27
  - rate limiting 4-26

iSAN routing
- assign mSAN_IDs 4-58
- description 4-24
- fabric autonomy 4-24
- iFCP protocol 4-22
- proxy Domain_ID 31 4-25
- routing domain 4-25

iSCSI protocol
- description 4-59
- initiators 4-59
- server consolidation 4-60
- storage consolidation 4-61
- targets 4-59

ISL
- bandwidth 3-22
- large fabric support 3-30
- limitations 3-20
- oversubscription 3-32
- path selection 3-25
- port fencing 1-36
- preventing oversubscription 3-33

J

jumbo frames
- description 1-24
- optimizing WAN use 4-56

L

large fabric
- fabric initialization 3-29
- fabric scalability 3-39
- high-bandwidth ISLs 3-30
- high-port count directors 3-30
- problems 4-3

laser transceiver
- description 5-4
- restrictions 5-5
- SFP transceiver 5-8
- transmission distance 5-4

latency
- dark fiber transport 4-46
- directors 1-7
- IP transport 4-48
- SONET/SDH transport 4-47
- WDM transport 4-47

LC duplex connector 5-7

LIM
- assigning BB_Credits 4-50
- Intrepid 10000 Director 5-2

load balancing 3-22

local area network
- comparison to WAN 4-36
- latency 4-36
- protocol stack 4-36
- reliability 4-37

logical port addressing 3-42

M

MAC address
- director or fabric switch 6-9
- Eclipse 1620 SAN Router 6-9
- Eclipse 2640 SAN Router 6-9
- management server 6-8

management server
- CHAP authentication 5-16
- default network address 6-8
- description 2-7
- Ethernet connectivity 5-10
- illustration 2-7
- installation planning 6-5
- minimum specifications 2-8
- plan security measures 6-11
- product overview 1-3
- recommended specifications 2-9

map, port card 3-43

mesh fabric
- description 3-14
- illustration 3-15
- suitability 3-16

mFCP protocol
- comparison to iFCP 4-28
- connecting SAN routers 4-20
- description 4-20

mSAN routing
- allocating Zone_IDs 4-30
- description 4-18
- mFCP protocol 4-20
- proxy Domain_ID 30 4-18
- routing domain 4-18
- supported limits 4-21

multimode cabling
- 50/125 5-7
- 62.5/125 5-7

multiswitch fabric topology
- best practices 3-20
- connecting FC-AL devices 3-11
- description 3-2
- domain ID assignment 3-24
- E_Port segmentation 3-27
- fabric initialization 3-29
- fabric performance 3-30
- fabric WWN assignment 3-24
- FCP and FICON intermix 3-41
- frame delivery order 3-26
- illustration 3-19
- ISL bandwidth 3-22
- load balancing 3-22
- multiple data transmission speeds 3-51
- performance objectives 3-20
- planning considerations 6-29
- preferred path 3-22
- principal switch selection 3-23
- topology limits
  - fabric elements 3-19
  - hop count 3-20
  - ISLs 3-20
  - vendor interoperability 3-20
- zoning configurations 3-28

N

N_Port DHCHAP authentication 5-17
N_Port ID virtualization feature 5-39
name conventions, ports 6-13
name server zoning
  introduction 1-28
  planning requirements 6-30
network addresses
  default settings 6-8
  planning 6-7
nickname conventions, ports 6-14
nonresilient fabric
  dual 3-38
  single 3-37
nScale architecture 1-12

O

open-system management server
  description 5-35
  introduction 2-5
  plan console support 6-6
OpenTrunking feature
  description 5-37
  planning considerations 3-22
  support planning 6-29
operating environment A-7
optional feature key
  description 5-33
  Element Manager application 5-39
  Flexport Technology 5-36
  FMS 5-35
  format 5-34
  full fabric 5-38
  full volatility 5-38
  N_Port ID virtualization 5-39
  OpenTrunking 5-37
  OSMS 5-35
  remote fabric 5-39
  SANtegrity Authentication 5-36
  SANtegrity Binding 5-36
OSMS feature
  description 5-35
  introduction 2-5
  plan console support 6-6
out-of-band product management
  command line interface 2-3
  EFCM 2-2
  EFCM Basic Edition interface 2-3
  EFCM Lite application 2-3
  Element Manager application 2-2
  SAN management application 2-2
  SANavigator 2-2
  SANvergence Manager 2-2
  SNMP 2-3
oversubscription, ISL 3-32

P

password
  authentication 5-16
  protection 1-31, 5-15
PCP user database 5-17
PDCM arrays
  description 5-20
  planning considerations 5-20
performance features
  directors 1-7
  fabric switches 1-15
  SAN routers 1-23
performance tuning a fabric 3-35
persistent binding 5-29
planning checklist
  operational setup tasks 6-36
  planning and hardware installation tasks 6-35
planning tasks
  assign port names and nicknames 6-13
  complete planning checklists 6-34
  complete planning worksheet 6-14
  consider interoperability with end devices 6-4
  diagram planned configuration 6-12
  establish security measures 6-11
  plan AC power 6-28
  plan console management support 6-5
  plan e-mail notification 6-11
  plan Ethernet access 6-7
  plan Fibre Channel cable routing 6-3
  plan multiswitch fabric 6-29
  plan network addresses 6-7
  plan phone connections 6-12

plan SAN routing 6-31
plan SNMP support 6-10
plan zone sets 6-30
prepare a site plan 6-2
planning worksheet 6-14
port
binding 1-31
blocking 1-28
fiber-optic cabling 5-1
logical port addressing 3-42
name conventions 6-13
nickname conventions 6-14
numbering 3-42
port card
Intrepid 6064 Director 5-2
Intrepid 6140 Director 5-2
map 3-43
port connections
any-to-any connectivity 1-28
Eclipse 1620 SAN Router 1-25
Eclipse 2640 SAN Router 1-27
Intrepid 10000 Director 1-13
Intrepid 6064 Director 1-9
Intrepid 6140 Director 1-11
port blocking 1-28
Sphereon 3232 Fabric Switch 1-16
Sphereon 4300 Fabric Switch 1-17
Sphereon 4400 Fabric Switch 1-19
Sphereon 4500 Fabric Switch 1-20
Sphereon 4700 Fabric Switch 1-22
zoning 1-28
port fencing (E_Port or ISL) 1-36
power requirements
directors A-3
fabric switches A-3
Fabricenter cabinet A-8
planning considerations 6-28
SAN routers A-3
preferred path
description 5-23
planning considerations 3-22
principal switch selection 3-23
private
arbitrated loop 3-9
loop device connectivity 3-7
product management
command line interface 2-3
EFCM 2-2
EFCM Basic Edition interface 2-3
EFCM Lite application 2-3
Element Manager application 2-2
FMS feature 2-5
inband methods 2-5
management interface summary 2-6
OSMS feature 2-5
out-of-band methods 2-2
SAN management application 2-2
SANavigator 2-2
SANvergence Manager 2-2
SNMP 2-3
product overview
connectivity features 1-28
directors 1-2
EFCM application 1-5
fabric switches 1-2
Fabricenter cabinet 1-3
management server 1-3
SAN routers 1-3
SANavigator application 1-5
SANvergence Manager application 1-5
security features 1-31
serviceability features 1-32
protocol intermix
best practices 3-47
fabric element management 3-41
impacting features 3-46
management limitations 3-44
port numbering and logical port addressing 3-42
protocols
Fibre Channel 3-41
FICON 3-41
iFCP 4-22
iSCSI 4-59
mFCP 4-20
proxy routing domains
Domain_ID 30 4-13, 4-18
Domain_ID 31 4-13, 4-25
iSAN routing 4-14
logical device connectivity 4-14
mSAN routing 4-14

public
arbitrated loop 3-8
loop device connectivity 3-6

R

R_Port
configuring 4-34
Domain_ID assignment 4-15
operation 4-11
RADIUS server support 5-18
rate limiting
description 4-26
implementation 4-51
intelligent port speed selections 4-52
recovery point objective
dark fiber transport 4-46
description 4-37
IP transport 4-48
SONET/SDH transport 4-47
WDM transport 4-47
recovery time objective
dark fiber transport 4-46
description 4-37
IP transport 4-48
SONET/SDH transport 4-47
WDM transport 4-47
redundant fabrics 3-38
remote data replication
asynchronous mode 4-38
dark fiber transport 4-46
IP transport 4-48
SONET/SDH transport 4-47
synchronous mode 4-38
WDM transport 4-47
remote fabric feature
assigning BB_Credits 4-50
description 5-39
Intrepid 10000 Director 1-29
remote workstations
description 2-10
Ethernet connectivity 5-11
installation planning 6-6
minimum specifications 2-10
resilient fabric
dual 3-38
single 3-38
restore features
CD-RW drive 2-13
director and fabric switch NV-RAM configuration 2-13
SAN management data directory 2-14
SAN router NV-RAM configuration 2-14
role-based FlexPars 4-7

S

SAN island
benefits 4-2
characteristics 3-18
consolidation
FlexPar technology 4-4
SAN routing 4-8
description 3-18
problems 4-2
SAN management application
EFCM application 1-5
main window (SAN router) 2-18
product overview 1-5
SANavigator application 1-5
SANvergence Manager application 1-5
SAN router
description 1-22
IRL optimization 4-25
logical connectivity 4-12
performance features 1-23
physical connectivity 4-11
product overview 1-3
proxy Domain_ID 30 4-13, 4-18
proxy Domain_ID 31 4-13, 4-25
router fabric manager 4-15
router name server 4-19
serviceability features 1-32
specifications A-1
zoning
append IPS zones 4-17
no zone synchronization 4-16
SAN routing
best practices 4-29
description 4-8
iFCP protocol 4-22
inter-FlexPar routing 4-28
iSAN routing 4-24
mFCP protocol 4-20

mSAN routing 4-18
planning requirements 6-31
R_Port operation 4-11
routing domain (iSAN) 4-25
routing domain (mSAN) 4-18
SAN island consolidation 4-8
Tier 1 (fabrics) 4-8
Tier 2 (mSANs) 4-9
Tier 3 (iSANs) 4-9
zone policy 4-16
SANavigator application
description 2-15
GUI description 2-15
product overview 1-5
SANtegrity Authentication feature
CHAP authentication 5-16
CT authentication 5-17
description 5-16
DHCHAP authentication 5-17
feature key description 5-36
inband access control list 5-18
out-of-band access control list 5-18
password safety 5-16
PCP user database 5-17
RADIUS server support 5-18
security log 5-18
SSH protocol 5-18
support planning 6-30
SANtegrity Binding feature
description 5-19
Enterprise Fabric mode 5-19
fabric binding 5-19
feature key description 5-36
planning considerations 5-20
support planning 6-30
switch binding 5-19
SANvergence Manager application
description 2-18
GUI description 2-18
main window 2-18
product overview 1-5
security provisions
best practices 5-30
general description 5-15
password protection 5-15
PDCM arrays 5-20
persistent binding 5-29
preferred path 5-23
SANtegrity Authentication 5-16
SANtegrity Binding 5-19
security feature description 1-31
security log description 5-18
server-level access control 5-29
storage-level access control 5-30
zoning 5-25
server
consolidation 3-12
EFCM Basic Edition interface connectivity 5-14
iSCSI server consolidation 4-60
server-level access control 5-29
SNMP connectivity 5-13
serviceability features 1-32
SFP optical transceiver
description 5-4
illustration 5-8
longwave laser 5-4
restrictions 5-5
shortwave laser 5-4
transmission distance 5-4
shared mode operation
description 3-3
illustration 3-3
shipping environment A-6
singlemode cabling, 9/125 5-7
site plan preparation 6-2
SNMP
management workstations 5-13
MIBs 5-13
product management 2-3
support planning 6-10
software
command line interface 2-21
EFCM application 2-15
EFCM Basic Edition interface 2-20
Element Manager application (director and fabric switch) 2-16
Element Manager application (SAN router) 2-19
SANavigator application 2-15
SANvergence Manager application 2-18

SONET/SDH distance extension
bandwidth 4-47
description 4-42
illustration 4-43
latency 4-47
recovery point objective 4-47
recovery time objective 4-47
specifications
director clearances A-5
director dimensions A-1
director heat dissipation A-5
fabric switch clearances A-5
fabric switch dimensions A-1
fabric switch heat dissipation A-5
fabric switch power requirements A-3
Fabricenter cabinet clearances A-8
Fabricenter cabinet dimensions A-8
Fabricenter cabinet footprint A-8
Fabricenter cabinet power requirements A-8
SAN router clearances A-5
SAN router dimensions A-1
SAN router heat dissipation A-5
Sphereon 3232 Fabric Switch
default network address 6-9
description 1-15
FRUs 1-15
illustration 1-15
Sphereon 4300 Fabric Switch
default network address 6-9
description 1-16
FL_Port connectivity 3-10
FRUs 1-17
illustration 1-16
Sphereon 4400 Fabric Switch
description 1-18
FL_Port connectivity 3-10
FRUs 1-18
illustration 1-18
Sphereon 4500 Fabric Switch
default network address 6-9
description 1-19
FL_Port connectivity 3-10
FRUs 1-20
illustration 1-19
Sphereon 4700 Fabric Switch
description 1-21
FL_Port connectivity 3-10
FRUs 1-21
illustration 1-21
SSH protocol description 5-18
state change notification 1-29, 3-28
storage environment A-6
storage-level access control 5-30
subnet mask
director or fabric switch 6-9
Eclipse 1620 SAN Router 6-9
Eclipse 2640 SAN Router 6-10
management server 6-8
switch binding 5-19
switched mode operation
description 3-4
illustration 3-4
synchronous remote data replication
dark fiber transport 4-46
description 4-38
latency limitations 4-38
WDM transport 4-47

T

tape device consolidation 3-13
telephone connection
call-home support 6-12
service support 6-12
Tier 1 fabric connections 3-17
Tier 2 fabric connections 3-18
Tier 3 fabric connections 3-18
topology
arbitrated loop topology 3-2
core-to-edge fabric 3-16
fabric topology limits
fabric elements 3-19
hop count 3-20
ISLs 3-20
vendor interoperability 3-20
mesh fabric 3-14
multiswitch fabric topology 3-2

V

vendor interoperability limitations 3-20

W

WDM distance extension
- bandwidth 4-47
- description 4-40
- illustration 4-41
- latency 4-47
- recovery point objective 4-47
- recovery time objective 4-47

weight
- directors A-1
- fabric switches A-1
- Fabricenter cabinet A-8
- SAN routers A-1

wide area network
- comparison to LAN 4-36
- dedicated bandwidth 4-55
- latency 4-36
- optimize use
  - buffering 4-55
  - data compression 4-55
  - FastWrite technology 4-56
  - flow control 4-55
  - jumbo frames 4-56
- protocol stack 4-36
- rate limiting 4-51
- reliability 4-37

worksheet, planning 6-14
WWN assignment, fabric 3-24

Z

zone FlexPars 4-6
zoned fabrics
- configuration rules 3-28
- joining 5-28
- SAN router
  - append IPS zones 4-17
  - no zone synchronization 4-16
- SAN routing environment 4-34
- zone policy (SAN routers) 4-16

zoning
- benefits 5-26
- configuring zones 5-26
- description 5-25
- introduction 1-28
- joining zoned fabrics 5-28
- planning considerations 5-28
- planning requirements 6-30
- SAN router
  - append IPS zones 4-17
  - no zone synchronization 4-16
- SAN routing environment 4-34
- zone policy (SAN routers) 4-16
- zone sets 5-27

zoning policy (SAN routers)
- append IPS zones 4-17
- implementing 4-34
- no zone synchronization 4-16